University of Alberta

3D Inter-Fractional Patient Set-Up Verification and the Biological Impact of Positioning Errors

by

Hans-Sonke Friedrich Jans

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements of the degree of *Master of Science*

in

Medical Physics

Department of Physics

Edmonton, Alberta

Spring 2006

Library and Archives Canada

Published Heritage Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et Archives Canada

Direction du Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence
ISBN: 0-494-13831-9
Our file Notre référence
ISBN: 0-494-13831-9

NOTICE:
The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:
L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

Canada

Abstract

Reproducible patient positioning during fractionated external beam radiation therapy is imperative to ensure matching between delivered and planned dose distributions. A 2D-3D image registration method to verify the patient's three-dimensional setup (rotations, translations) using orthogonal portal images and megavoltage digitally reconstructed radiographs (MDRRs) is presented. Registration precision improves with additional image pre-processing and interpolation of the parameter space. For combined rotations and translations it is better than 0.5 mm and 0.3° (1 SD) if translation $\leq$7 mm and rotation $\leq$4°. Dependence on angular spacing between MDRRs, three cost functions, pixel size and field-of-view is investigated. A model for rotation-translation coupling is presented that mostly agrees with experimental data. The registration method is applied to two clinical prostate treatment plans. The patients' dose distributions, dose-volume-histograms and changes to tumor control and normal tissue complication probabilities (TCP/NTCP) are calculated. TCP remains constant, whereas patients' rectal NTCP increases due to patient positioning errors.

Acknowledgements

This work was possible only with the help and support from a number of people for whose contributions I am truly grateful. I would like to thank my supervisor Dr. Gino Fallone, whose comprehensive knowledge of and tangible enthusiasm for the field of Medical Physics is an inspiration and motivation. He also was a great support when applying for scholarships and provided the opportunity to work on very interesting, hands-on quality assurance projects. I also greatly respect Dr. Satyapal Rathee for his expertise and thank him for direction and encouragement at crucial junctures as well as practical help in this project, such as acquiring CT scans, the introduction to the Helax TPS, pointing out relevant literature or painstakingly revising manuscripts. I would also like to acknowledge the efforts of Dr. Mauricio Sacci and Dr. Richard Sydora for their thoughtful consideration and review of my thesis.

Everyone in the Medical Physics department has been exceptionally helpful and approachable and made this a tremendous place to study and work. At the risk of forgetting someone, I would like to acknowledge: Dr. Alasdair Syme for familiarizing me with the operation of linacs, EPIDs and the Pinnacle TPS. He and also Dr. Stephen Steciw are great guys to discuss ideas with, as well as being a source of encouragement at rough spots during the course of the project. Brad Murray provided patient images, Varis connectivity and a drive home for me late at night. Sherry Connors did not mind after-hours calls for help with linac interlocks. UNIX and Windows machines were expertly trouble-shot by Lee Santon, Ernie Mah and Mike McDonnell. Debbie, Maureen and Jocelyn went the extra mile whenever administrative issues needed to be tackled.

Contributions from outside the Medical Physics department included Dr. Todd McNutt and Karl Bzdusek from Philips Medical Systems, Madison, WI, for help in writing Pinnacle plug-ins and Dr. Philippe Thévenaz (L'Ecole Polytechnique Fédérale de Lausanne, Switzerland) who assisted in compiling his 2D registration code on my work station. I am also grateful for a scholarship received from the Alberta Heritage Foundation for Medical Research.

Two then-graduate students confirmed my interest in Medical Physics by sharing their experiences and enthusiasm: Dr. Geetha Menon and Dr. Brad Warkentin – thank you for taking the time; I knew then that I was on to something.

The two semesters of lectures were a great preparation for the thesis work and very enjoyable, largely due to excellent teaching. Dr. Jack Cunningham's humor and tales from the past made the field come alive; Dr. Gino Fallone provided a lively introduction that made sense of what could otherwise have been dry imaging theory; Dr. Ron Sloboda's thorough and systematic teaching was a pleasure to follow; Dr. Don Robinson was able to liven up weeks of repetitive multiplication (sometimes by recruiting extraterrestrial help); Dr. Terry Riauka and Dr. Satyapal Rathee pulled the world of applied medical imaging together by discussing real-life applications; Dr. Rick Hooper explained nuclear medicine the way he wished it had been taught to him and Dr. Dave Murray's broad sweep of cell biology was an unforgettable excursion into the miracle of life.

I would like to thank my fellow students at the Cross Cancer Institute, those with whom I started the program and who were a great team to study and go through the classes with: Amr, Colleen, Deluan, Laura, Lesley, Matt, Ryan and Sandra; as well as those who began their studies in Medical Physics either earlier or later: Charlie, David, Donata, Erin, Michael, Steven, Tara, Teo, Yingli – thank you for good talks, fun and friendship.

Last but not least I would like to thank my extended family, both in Germany and Canada, especially my parents, Hans-Uwe and Heike and parents-in-law, Jack and Sue – your love and support for the Jansfamily during these years were invaluable.

Finally, to those who matter most to me: my wonderful wife Sharon and children Eliesa, Christian and David. Your love and standing by my side through thick and thin blesses my life and work immeasurably and has helped me in countless ways. I love you more than I can say and could not be prouder of my family. Soli Deo gloria.

To

Sharon

Eliesa, Christian, David

Table of Contents

Chapter 1: Introduction **1**

1.1. Scope and Organization of this thesis 1

1.2. Overview 3

1.3. Patient positioning in external beam radiation therapy 5

1.4. Biological impact of patient positioning errors 9

1.5 References 11

Chapter 2: Image Registration: Materials and Methods **15**

2.1. Overview 15

2.2. Hard- and Software 18

2.3. Generation of the MDRR data base 19

2.4. Acquisition of the EPID Image Pairs 22

2.5. Image pre-processing 24

2.6. Cost Functions 30

2.7. Registering 3D Translations 32

2.7.1 Searching the MDRR data base 32

2.7.2 2D-translations 33

2.7.3 3D-translations 34

2.8. Registering 3D Rotations 35

2.9. Coupling of Rotations and Translations 37

2.10. Testing the Method 41

2.10.1 Translations 42

2.10.2 Rotations 43

2.10.3 Coupling of Rotations and Translations 45

2.10.4 Combined Translations and Rotations 46

2.11. References 48

Chapter 3: Image Registration: Results and Discussion **49**

3.1. Evaluation Overview 49

3.2. Translations 51

3.3. Rotations 53

3.4. Coupling of Rotations and Translations 61

3.5. Combined Translations and Rotations 63
3.6. Summary 65
3.7. References 66

Chapter 4: Biological Impact: Materials and Methods 67
4.1. Overview 67
4.2. Patient Position 69
4.3. Clinical Cases 69
4.4. Dose Calculation 71
4.5. Radiobiological Model 73
4.5.1. TCP model 73
4.5.2. NTCP model 75
4.6. References 77

Chapter 5: Biological Impact: Results and Discussion 78
5.1. Scope 78
5.2. Setup errors 78
5.3. Dose Distribution and Dose-Volume Histograms 80
5.4. TCP and NTCP 82
5.5. Summary 84
5.6. References 85

Chapter 6: Summary and Conclusion 86
6.1. Image Registration 86
6.2. Biological Impact 89
6.3. Future Developments 89
6.4. References 92

Appendix 93
A.1. Patient and Linac rotation angles 93
A.1.1 Coordinate transform 94
A.1.2 Extracting linac angles from the transformation matrix M 96
A.2. Standard Deviation of a Uniform Random Distribution 99
A.3. References 100

Bibliography 101

List of Tables

Table 2.1 Translations along any of the three orthogonal axes appear as rotations about axes orthogonal to the beam (projection) axis and to the direction of translation. Because orthogonal images are acquired (one in AP and one in LR view), translations in SI direction are shown in both, the AP and LR view image..... 41

Table 3.1 Registering known random translations: Values describe the distribution of differences between setup and found (registered) translations (mm). They are the averages of values obtained for registrations using each of the three cost functions..... 52

Table 3.2 Registering known random rotations: Values represent the differences between angles found by the registration routine and known setup angles (deg). Values are for extended image pre-processing (bottom row of charts in Figure 3.3).. 54

Table 3.3 Registration performance as a function of MDRR angular interval for the "Parabolic fit" strategy (Windows XP workstation equipped with a 2 GHz Xeon processor and 1 GB RAM). 58

Table 3.4 Registration performance for combined random translations and rotations. Random rotations were $\leq 4°$ about each axis; random translations were within a ± 15 mm (top rows) and ± 7 mm (bottom rows) interval. 64

Table 4.1 Treatment parameters for the two patients 70

Table 5.1 Averages, standard and maximum deviations for clinical patient setup errors, derived from the measurement of fractions per patient shown in Figure 5.1..... 81

Table 5.2 Changes of NTCP and TCP as a result of setup error 83

List of Figures

Figure 1.1 Example of a modern seven field head-neck IMRT treatment plan. The isocentric slice (right) shows the dose distribution conforming tightly to a C-shaped PTV, while sparing the spinal cord. 4

Figure 1.2 The Clinical Target Volume (CTV) is expanded to the Planning Tumor Volume (PTV) in order to account for possible positioning errors. This expansion can cause increased irradiation of an Organ At Risk (OAR), located close to the GTV. 6

Figure 2.1 Flowchart outlining the 3D image registration method employed in this thesis 16

Figure 2.2 The patient coordinate system; translations along and rotations about each of the three orthogonal axes are registered with the method describe in this chapter. 17

Figure 2.3 The humanoid phantom ("Pixy"). 19

Figure 2.4 Pinnacle's scaling parameter for attenuation coefficients was set to achieve best visual agreement between MDRR image (left and middle) and EPID image (right). Both MDRRs were generated from the same CT data set, left scaled for low and right for high (MeV range) energy. 20

Figure 2.5 MDRR images are generated by ray-tracing through the CT data set, which is depicted in the center (here part of the torso of the humanoid phantom). AP images are calculated with simulated beams that originate from a source above the CT data set and are propagated trough the CT data set from top to bottom. LR images are the result of simulated beams traversing the CT data set from right to left. 21

Figure 2.6 Acquisition of orthogonal EPID images of patients; AP view (left) and LR view (right). 23

Figure 2.7 Equivalency of phantom roll around SI axis rotation (left) to the adjusted gantry angle (right) that is used to emulate rotations of the humanoid pixie phantom. The phantom rotations about the other axes wre similarly realized using combinations of couch, gantry and collimator angles of the linear accelerator. 24

Figure 2.8 The EPID image (top) and the MDRR image (bottom) are cropped and re-sampled to the same pixel size; shown here is only the AP view. 25

Figure 2.9 Image pre-processing steps used in the method. The original (raw), cropped EPID and MDRR images (first column); Original images convolved with a large Gaussian kernel (second column); Blurred images subtracted from original images (third column); Histograms matched to a normal distribution (fourth column); Histograms aligned to same interval (fifth column). After unsharp masking overall image intensities are more balanced and finer structures (e.g. vertebral bodies) become visible.

Histogram matching enhances contrast and further enhances similarity. Finally the image histograms are shifted into the same gray-value range of [0, 255]; pixel values are non-integer (double). For consistent visualization, all images shown have been scaled to gray values between 0 and 255..... 26

Figure 2.10 The position of the beam axis on the EPID image is determined by the cross shaped pattern generated by the graticule filter (left). It is found by calculating the center of mass (CM) coordinates of the edge pixels of the central dot, visible in the magnified image on the right. 27

Figure 2.11 Epid image of a 10 x 10 cm2 field (left). The magnified image on the right shows the edge pixels of the square field (black). The beam axis on the EPID image was found by calculating the CM of those edge pixels. The images here are displayed in false color for increased visual contrast..... 28

Figure 2.12 Rotational offset of the EPID images was found by exposing a square field. Initially, the edge pixels (black) do not form a straight line (top images, marked by arrows). After rotating the image by 0.35°, the edge pixels form a straight line. The images here are displayed in false color for increased visual contrast. 29

Figure 2.13 Part of the pre-processing of clinical EPID images is a routine that automatically removes the cross pattern of white dots (left image), which are caused by the gantry-mounted graticule filter. The routine replaces pixel values within the dots with local averages around them (right image). 30

Figure 2.14 The 2D registration is performed separately for LR and AP view images. The AP view registers translations in S-I and L-R direction, whereas the LR view is sensitive to translations in S-I and A-P directions. The two results obtained from each of the views for the SI translations is averaged. 34

Figure 2.15 Two-dimensional cross section of the angle parameter space. Left the full parameter space, right the values around the minimum, to which a parabola is fitted. For display purposes, the parabola has been elevated along the vertical axis above the surface of the cost function..... 36

Figure 2.16 Geometry illustrating how the translation of an object (b) can appear similar to a rotation (c). Originally, both details A and B within the square object are aligned such that they over lap in the imaging plane. The details are separated in the case of both, translation (b) and rotation (c). SID = Source-to-Image-Distance, SOD = Source-to-Object-Distance..... 38

Figure 2. 17 Pseudo rotation angles as a function of the translation of an object, plotted for different SOD values according to Equation 2.11. 40

Figure 3.1 Flow chart illustrating the how the performance of the registration method is evaluated by comparing registered (found) with setup values, and characterizing the distribution of differences by their average, standard deviation and maximum deviation, corresponding to the dots, bars and whiskers respectively in the bottom chart. 50

Figure 3.2 Registration results for translations, investigating the effect of using different cost functions and image pre-processing. The charts are evaluating the difference between the translation value found by the algorithm and the random setup values along the AP, LR and SI axes. The average, standard and maximum deviations of the differences are shown as dots, thick error bars and thin error bars respectively. 51

Figure 3.3 Effect of image pre-processing and cost function on registration of rotation angles about each axis. Evaluated are the difference between the angles found by the registration routine and the random, known setup value. Shown is the average value (dots), standard deviation (thick bar) and min and max values (thin bars). Averages close to zero and small bars indicate accurate and precise registration. 52

Figure 3.4 Angles found by the registration routine (measured angle) vs. the known setup angles by which the phantom was rotated. Shown are results for each of the three axes (AP, SI, LR) for the cost function based on mutual information. The data has been fitted with a straight line. 55

Figure 3.5 Dependence of registration results on angular interval of MDRR data base: Charts evaluate the difference between the setup and found angles of the phantom using three different strategies: "MDRR angle", "Parabolic fit" and "From translations" (each graph's x-axis has been offset slightly to avoid overlap of the data). The left chart shows the average difference, which should be zero for ideal registration. To avoid clutter, values have been averaged for the different cost functions, as well as axes; the 'error bars' depict the range of those values. The middle chart shows the standard deviation of the angle difference. The 'error bars' again indicate the range of values for the different axes and cost functions. The chart on the right shows the maximum deviation of the angle difference from its average value (for all three cost functions and all three axes); it depicts a 'worst case' scenario. 56

Figure 3.6 Histogram showing the distribution of misregistration (difference between the angle values found and the setup value) for the three different search strategies at an MDRR interval spacing of 2.5°, for all three axes. 57

Figure 3.7 Different FOV tested with the routine. Their sizes varied from 5x5 cm2 to 14x14 cm2. Shown are the AP (left) and LR (right) view EPID images of the humanoid phantom after pre-processing. 59

Figure 3.8 Dependence of registration results on the size of the FOV: Charts evaluate the angle difference between the setup and found angles of the phantom. To reduce clutter in the chart, values have been averaged for the three different axes; the 'error bars' depict the range of these values. 59

Figure 3.9 Dependence of registration results on the number of pixels in a constant 10 x 10 cm2 FOV. Charts evaluate the angle difference between the setup and found angles of the phantom. To reduce clutter in the chart, values have

been averaged for the three different axes; the 'error bars' depict the range of these values. 60

Figure 3.10 Pseudo rotation angles found by the registration method: These occur when the object (in this case the humanoid phantom) is translated because the translated projection image through the object appears similar to an image obtained when the same object has been rotated. 61

Figure 3.11 Registered values for translations as a function of rotations about the AP, SI and LR axes (left, middle and right chart, respectively) for the un-translated phantom. Error bars for each axis are the averages of values obtained for phantom and electronic translations in Table 1. Each graph's x-axis has been offset slightly to avoid overlap of the data. 63

Figure 4.1 Flow chart outlining how patient setup errors (translations and rotations) were processed to evaluate their biological impact on the patient, expressed as TCP/NTCP. 69

Figure 4.2 BEV of a main field (left) and coned down boost field (right). The beam aperture is shaped with multi-leaf-collimators (MLC, horizontal stripes), which block the beam. The aperture is shaped by the MLC to conform to the shape of the PTV, which contains the prostate (yellow or gray). The width of each MLC leave is 0.5 cm and the height of the PTV is approx. 8.5 cm. The darker (green) structure running along the right side is caused by overlap of the prostate and the rectum (blue). 70

Figure 4.3 Transverse isocenter slice of a "master" treatment plan after it has been set up on Pinnacle. Seven beams were used in this plan. 72

Figure 4.4 Plot of the curves for TCP (Equation 4.6) and NTCP curve (Equation 4.7) for the prostate and rectum, respectively, using the parameters cited in the text. 75

Figure 5.1 Measured setup errors for patient A (top) and B (bottom) over the course of their treatment as indicated by the fraction number on the x-axis 80

Figure 5.2 Isocentric transverse slice of a prostate treatment plan for patient A, showing the PTV, rectum and isodose lines for the original plan (left) and the summed dose for all setup errors (sum-plan, right). Higher % isodose lines in the sum-plan have moved towards the rectum as a result of setup errors. 100% corresponds to a dose of 72.9 Gy. 81

Figure 5.3 Difference of the dose distributions of the sum-plan and the original plan of patient A for the isocentric slice shown in Figure 5.2. An overdose by 250 cGy for large parts of the rectum and a small portion of the prostate is evident. 82

Figure 5.4 Dose-volume-histograms for the PTV and rectum for the original plan and the sum-plan (containing the setup errors) for patient A (left) and B (right). 82

Figure A.1 Coordinate transforms between the Patient, BEV and Room coordinate systems. 95

List of Abbreviations

ART	Adaptive Radiation Therapy
BEV	Beams Eye View
CM	Center of Mass
CT	Computed Tomography
CTV	Clinical Target Volume
DICOM	Digital Imaging and Communications in Medicine
Dof	Degrees of freedom
DRR	Digitally Reconstructed Radiograph
EBRT	External Beam Radiation Therapy
EPID	Electronic Portal Imaging Device
EUD	Equivalent Mean Dose
FOV	Field Of View
IMRT	Intensity Modulated Radiation Therapy
Linac	Linear accelerator
MDRR	Megavoltage DRR
MLC	Multi-Leaf Collimator
MU	Monitor Unit
NTCP	Normal Tissue Complication Probability
OAR	Organ At Risk
PTV	Planning Target Volume
RAM	Random Access Memory
ROI	Region Of Interest
TCP	Tumor Control Probability
TPS	Treatment Planning System
SDR	Sigmoidal Dose Response
SID	Source-to-Image Distance
SOD	Source-to-Object Distance
VOI	Volume Of Interest

Chapter 1

Introduction

1.1. Scope and Organization of this thesis

Modern delivery of external beam radiation therapy (EBRT) requires precise positioning of the patient in the treatment beam. Failure to do so can compromise the treatment outcome by either delivering insufficient radiation to the tumor or heightened dose levels to the surrounding healthy tissue. This thesis is concerned with those two aspects of treatment delivery: patient positioning and the biological effects of positioning errors.

The first goal of this thesis is to develop and validate a method for the accurate measurement of a patient's position relative to the treatment beams. The method should use readily available equipment, such as a treatment planning system (TPS) and an electronic portal imaging device (EPID) and require minimum user interaction. Information about patient position can then be derived by comparing EPID images of the patient on the treatment couch with projections through the patient's CT data set, calculated by the TPS. The patient position should be measured in 3-dimensions (translations and rotations about three orthogonal axes), i.e. 6 degrees of freedom (dof). Non-rigid transformation of patient anatomy is not considered because EPID images are

projections through the patient that do not provide 3D anatomic information and also lack sufficient soft tissue contrast as a result of the high photon energy with which they are acquired (the 6 MV electron acceleration potential of the linear accelerator results in a bremsstrahlung spectrum with an average photon energy of approx. 2 MeV). Therefore, internal patient anatomy is not sufficiently delineated and changes in internal anatomy cannot be registered.

The second goal of this thesis is to use the information thus acquired about patient positioning by evaluating its impact on treatment outcome, i.e. the biological effect of positioning errors. This is done by calculating, in retrospect, the dose distribution delivered to the patient during the therapy with the true (measured) patient position, i.e. including setup errors. This dose distribution provides a more realistic picture of the treatment and allows a truer judgment of the treatment outcome and the severity of the setup errors. The biological impact of the patient positioning errors is quantified by calculating changes to the probability for tumor control (Tumor Control Probability, TCP) and normal tissue complication (Normal Tissue Complication Probability, NTCP).

The information about patient position can also potentially be used to correct patient setup errors in real time, while the patient is positioned on the treatment couch and before the treatment dose is delivered. Implementation of this procedure is beyond the scope of this thesis as it requires the change of treatment protocols, changes to clinical hard- and software as well as a further increase in processing speed of our registration algorithm. However, the theoretical description of how to geometrically account for patient setup errors in radiation delivery is presented.

In correspondence with the two goals mentioned above, this thesis is structured into two parts: Chapter 2 describes the methods and materials used for measuring patient position by utilizing a pair of orthogonal megavoltage radiographs (EPDI images). Chapter 3 presents the results regarding the accuracy of the method and its sensitivity to various steps and parameters. The registration parameters resulting in the most accurate registration are then used to register clinical patient data. In the second part of this thesis, the results of these registrations are in turn evaluated for their

influence on dose distribution, TCP and NTCP: chapter 4 describes the materials and methods used for this purpose and chapter 5 illustrates the results obtained. Finally, chapter 6 summarizes the conclusion of this thesis and provides an outlook for possible future projects building on the results. Appendix A.1 presents the transformation of coordinates between the patient and collimator coordinate systems, which are used several times throughout this work. Appendix A.2 calculates the standard deviation of a uniform distribution, used in chapter 3 to compare different registration strategies for 3D rotations.

Parts of the contents of this thesis have been accepted for publication in the "Medical Physics" journal. [1]

1.2. Overview

EBRT, used either by itself or in conjunction with surgery and/or chemotherapy, is one of the primary treatment modalities for most types of localized cancers. When delivered in sufficient intensity, ionizing radiation destroys the biological tissue it traverses. The energy deposited by the radiation – mostly in form of electronic ionization – causes death, growth arrest or malfunction of the cells forming the tissue. While those effects are desirable for cancerous tissue, they are to be minimized for normal tissue surrounding the target region. However, exposure of normal tissue to the treatment beam cannot be fully avoided. The tolerance of normal tissue to radiation dose (before occurrence of serious complications) therefore places limits on the amount of radiation dose that can be delivered to the cancer site. Not all normal tissue displays the same sensitivity to ionizing radiation. Radiation treatment plans are designed to minimize the dose delivered by the treatment beams to the most radiation sensitive tissues (critical structures). Geometric accuracy of dose delivery thus is crucially important for the success of the radiation therapy: it ensures accurate delivery of dose to the target region and in cases where critical structures are located in close vicinity to the tumor, their accurate avoidance. Missing either one of those goals, by either not delivering enough dose to parts of the tumor or by delivering excess dose to parts of a critical structure would render the treatment a failure; in the first case

because the likelihood of controlling the cancer would be compromised; in the second case because radiation-induced complications could severely limit the patient's quality of life after the treatment.

Two problems need to be solved in order to maximize dose to the tumor and simultaneously minimize dose to critical structures. First is the ability to shape the dose distribution with adequate precision to conform tightly to the target region. To this end, radiation therapy has seen the continuous development of planning and delivery techniques. Three-dimensional conformal radiation therapy (3D-CRT) and intensity modulated radiation therapy (IMRT)[2] represent successful efforts focused on delivering radiation dose to the tumor site in an ever more precise manner both in terms of the absolute dose and the geometrical accuracy with which the dose is 'painted' onto the target volume (Figure 1.1).

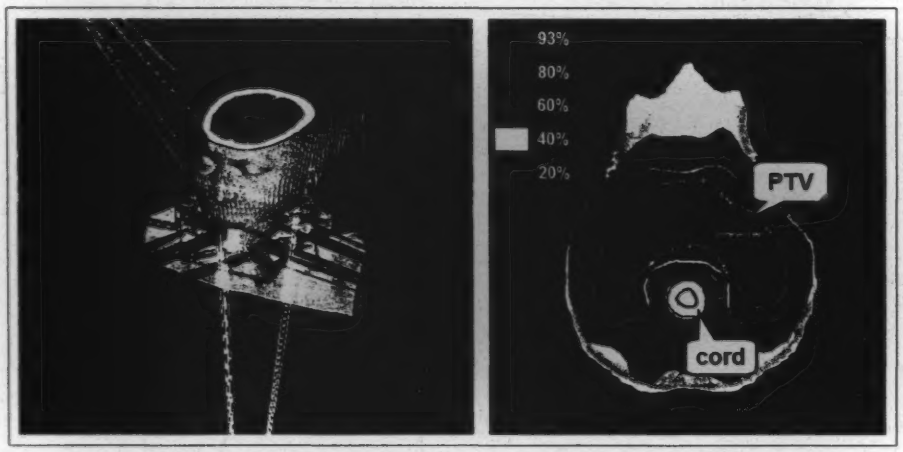

Figure 1.1 Example of a modern seven field head-neck IMRT treatment plan. The isocentric slice (right) shows the dose distribution conforming tightly to a C-shaped PTV, while sparing the spinal cord.

Secondly, the patient must be positioned correctly relative to the radiation beams delivering the treatment. Radiation treatment is delivered in multiple fractions, typically between 20 and 40, one fraction being delivered each weekday. Each time, the patient needs to be placed on the treatment couch in a consistently accurate way. The measurement of patient position presents a problem that has not been solved in a generally accepted, satisfactory manner and constitutes one of the main aspects of this

work. Several approaches taken to measure patient position are briefly reviewed below, prior to the introduction of the method chosen for this thesis.

Both, the patient position and the dose distribution within the patient are specified with the help of a treatment planning system (TPS), consisting of a computer workstation with the ability to define the position of treatment beams relative to the patient and to calculate radiation dose delivered by the beams. In the treatment plans that are generated with the TPS, the patient anatomy is generally represented by a CT data set. Beam positions are defined relative to this CT data set and dose calculations are based on the electron density values stored in its 3D matrix of voxels. The CT data set therefore also serves as a reference for patient positioning. It is acquired only once prior to planning and delivery of the treatment.

1.3. Patient positioning in external beam radiation therapy

At the beginning of each treatment fraction, before dose is delivered, the patient needs to be positioned on the treatment couch, ideally in the same pose he or she held in the CT scanner at the time the planning CT data set was acquired. This would insure that the treatment beam positions with respect to the patient remain as they were planned in the TPS. Any patient positioning error will cause the actual dose distribution delivered in a given treatment fraction to deviate from the planned one. Great care is therefore taken when positioning the patient on the treatment couch in order to reproduce his or her position at the time of the CT scan. This process is aided by means of fiducial (skin) markers that were placed on the patient when the CT scan was acquired and are visible on the CT scan. These skin markers are aligned with orthogonal lasers that mark the isocenter in the treatment room. This procedure attempts to place the treatment isocenter, defined by the CT data set in the TPS, at the isocenter of the tele-therapy unit.

It is becoming more and more common to verify the patient's position by acquiring, at a dose that is relatively small compared to a treatment dose, a megavoltage radiographic image of the patient in treatment position using a linear accelerator (linac) beam and an EPID. The EPID image shows the patient's actual position and can be

compared with a digitally reconstructed radiograph (DRR) image. The DRR image is calculated by projecting, computationally, rays with the same geometry as the treatment beam through the CT data set (ray tracing) used for treatment planning. Visual comparison of EPID image and DRR is used to adjust the patient position by translating the treatment couch appropriately. However, those translations are limited to within the plane perpendicular to the incident beam direction as defined by the EPID image; a patient's rotational offset is not corrected. Thus even with these precautions, setup errors in the order of several mm are known to occur[3-6]. They are accounted for in the treatment planning process by introducing generally accepted margins, which increase irradiated area from the Clinical Target Volume (CTV) containing the tumor, to the Planning-Target-Volume (PTV). For example, margins of the prostate CTV are usually taken to be 5 mm. This method attempts to insure that the entire tumor receives the prescribed dose, even in the case of setup errors. However it also causes increased irradiation of the normal tissue surrounding the tumor and is especially disadvantageous in cases where critical structures, also referred to as organs at risk (OAR), are located close to the CTV (Figure 1.2). The larger margins will then cause an undesired increase in the dose delivered to the OAR and may lead to increased NTCP values for the OAR.

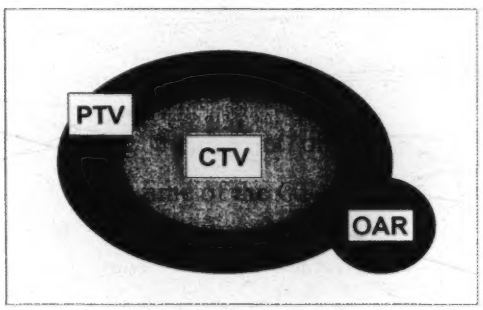

Figure 1.2 The Clinical Target Volume (CTV) is expanded to the Planning Tumor Volume (PTV) in order to account for possible positioning errors. This expansion can cause increased irradiation of an Organ At Risk (OAR), located close to the GTV.

Two factors can cause a change in the position of the CTV: an error in the setup position of the patient on the treatment couch or a shift or deformation of the CTV itself within the patient. In this thesis, the first cause is considered, i.e. the patient is considered to be a rigid body.

Prerequisite to correcting a positioning error is its measurement. Numerous approaches have been taken to improve the measurement of patient position with respect to the treatment beam during fractionated radiation therapy.

In a commercial system, several adhesive surface markers are placed close to the treatment site on the patient's skin. These markers are visible on the CT scan as well as with infra red cameras which are mounted in the treatment room. Semi automatic patient alignment is carried out by translating the treatment couch. This system achieves a precision between 2.0 and 3.5 mm (one standard deviation (SD)) [7].

The most advanced systems to date use a CT scanner mounted in the treatment room, dedicated solely to patient positioning. A CT scan is acquired prior to each treatment fraction while the patient is positioned on the treatment couch. After the CT scan is performed, the couch is turned 180 degrees around a central, vertical axis and placed underneath the linac. Since the patient remains motionless relative to the couch for the whole process, this approach significantly improves the registration precision to better than 0.5 mm (Court *et al.*[8]) or 0.16° and 0.16 mm (Kress *et al.*[9]). This system has the added advantage of visualizing the patient's internal soft tissue anatomy, potentially allowing adjustment of the treatment for internal organ motion and/or deformation. However, in-room CT imaging is not available on the majority of present clinical treatment installations.

The same is true for gantry mounted or fixed x-ray units in the treatment room. Those systems are used to acquire orthogonal views of the patient and achieve better bone-tissue contrast than images acquired with the linac's beam because of their lower effective photon energy. Various configurations of those kilovoltage imagers resulted in registration accuracies of 1.36 ± 0.11 mm (Ryu *et al.*[10]), better than 0.69 mm (1 SD, Pisani *et al.*[11]), 1.0 mm and 1.3° (1 SD, Murphy[12]), 1 mm (Lemieux *et al.*[13]) or 4.2 +/- 1.6 mm and 1.0 +/- 0.7° (Birkfellner *et al.*[14]). Recently, Fu *et al.*[15] published very precise results obtained using the kV-imagers of the commercial CyberKnife ® system: 0.3 mm for translations and 0.3° for rotations (rms errors).

Using one full rotation of the gantry, it is possible to acquire cone-beam CT images of the patient in treatment position. When acquiring the CT scan with linac's

MV beam[16], delineation of high-contrast tissue (bone, air, fat/water) would likely allow registration of rigid structures (bones), but results have not yet been reported in the literature.[17] A kV x-ray source-imager combination mounted on an additional c-arm on the gantry[18] allowed registration accuracy for the prostate of better than 2.0 mm along any axis.[19]

EPID imagers are most commonly utilized to determine patient position because of their availability (they are relatively inexpensive and installed on the majority of modern linacs) and ease of use.[20] However, since the images are produced at MV energies, their inherent contrast is lower than it is in kV imaging systems.[21] One strategy of using EPID imagers is to visualize markers implanted in the GTV. This allows for fast, accurate registration (± 0.5° [22]) and also account for organ motion. However, implanting the markers is invasive and it cannot be carried out in all treatment sites, marker delineation is usually a manual process and maker migration has been observed.[23]

In most clinical applications, the gray-value content of EPID images is registered to digitally reconstructed radiographs (DRRs).[20,24-30] Though in most cases only translations and in-plane rotations are measured, several studies have demonstrated the possibility to determine the 3D patient position (i.e. all 6 degrees of freedom (dof) that characterize a (rigid) body's pose) without the need for manual intervention or delineation of features when using pairs of EPID images.[31-33] This task is complicated by the decreased image contrast in MV images. Gilhuijs *et al.*[32] reported an accuracy of 1 mm for translations and 1° for rotations, Bansal *et al.*[31] found an accuracy of better than 1° for rotations and 2 to 3 pixels for translations. Sirois *et al.*[33] achieved standard deviations of 0.5 mm (less than one pixel) for translations and between 1.5° and 2.1° for rotations by registering orthogonal EPID images to Megavoltage DRRs (MDRRs).

In this thesis, the method by Sirois *et al.*[33] is expanded on, since it uses the readily available EPID imagers, does not require image segmentation (and thus minimizes the need for user intervention) and registers translations with good precision. Special attention was paid to improving the ability to register rotations. To this end, two additional steps were introduced: image pre-processing (which is shown to also have a

positive effect on the registration of translations) and a parabolic fit which interpolates the parameter space of cost functions. This approach was validated by registering several poses of a humanoid phantom that had been set up with known translations and rotations and evaluating the difference between the registered and the setup values.

1.4. Biological impact of patient positioning errors

As mentioned above, patient setup errors will alter the distribution of dose delivered to the patient and thus the biological response of both the tumor and normal tissue. Determining the change in the biological response constitutes the second main aspect of this thesis.

The impact of radiation on cancer can be characterized by the TCP, defined as the probability of local control of a tumor, and NTCP, defined as the probability of normal tissue complication associated with the treatment, usually assessed within a period of five years after treatment.[34] Calculating values for TCP and NTCP is a field of active research and various mathematical models have been suggested in order to predict and/or explain the values observed in clinical practice. The calculation of TCP for a tumor and NTCP for an organ is typically based on its dose-volume-histogram (DVH). A DVH is associated with a user-defined volume of interest (VOI) within the CT data set. A treatment plan typically contains VOIs for the CTV, PTV and for each critical structure affected by radiation. For a particular VOI, a DVH displays the fraction of volume of this VOI that has received a certain dose as a function of that dose. For the purpose of this thesis, a simple model based on Poisson statistics is chosen for the TCP and the phenomenological Sigmoidal Dose Response (SDR) model is used for the NTCP.[35]

Several approaches have been taken to study the effect of patient positioning errors. In a phantom study, Karger *et al.*[36] found that the surface dose in a spinal cord IMRT plan changes by 7.5% if a setup error of only 1 mm occurs, but the biological effect in terms of TCP and NTCP is not discussed. The effects of setup uncertainties on DVH have been studied by Cho *et al.*[37], who determined the changes in the DVH of the PTV in a lung treatment plan when introducing computer simulated setup errors. They

found significant changes in the DVH, but did not calculate TCP of the PTV. They also found that the dose distribution for a shifted patient position should be re-calculated for the new position rather than just shifting the previously calculated dose distribution by the amount of the patient setup error. In a patient study of prostate treatments, Song *et al.*[38] were able to boost the TCP of prostate cancer by between 1.2 % and 7.5 % while maintaining the same rectal NTCP when correcting for setup errors.

If the patient is not positioned on the treatment couch in the same way as in the CT scan (i.e. a positioning error has occurred), the treatment beams will "see" a different scenario and the dose distribution within the patient will change. In this work, a method is presented to evaluate the effect of patient positioning errors on the TCP and NTCP values in a few straight-forward steps:

- The 3D setup error is determined for each treatment fraction
- The dose distribution (including setup error) for each fraction is re-calculated
- The dose distributions for all fractions are added
- The DVHs are calculated from the added dose distributions
- The TCP and NTCP are calculated from the DVH

The values for TCP and NTCP are then compared with the ones obtained for the original treatment plan (i.e. no setup error has occurred). Their difference is a measure for the biological impact of the setup errors.

1.5. References

1 H.-S. Jans, A. Syme, S. Rathee, and B. G. Fallone, "3d inter-fractional patient position verification using 2d-3d registration of orthogonal images", Med. Phys., in press.

2 F. M. Kahn, "*The physics of radiation therapy*", 3 ed, (Lippincott Williams & Wilkins, 2003).

3 J. Hanley, M. A. Lumley, G. S. Mageras, J. Sun, M. J. Zelefsky, S. A. Leibel, Z. Fuks, and G. J. Kutcher, "Measurement of patient positioning errors in three-dimensional conformal radiotherapy of the prostate", Int. J. Radiat. Oncol. Biol. Phys. **37**, 435-444 (1997).

4 P. Remeijer, E. Geerlof, L. Ploeger, K. Gilhuijs, M. van Herk, and J. V. Lebesque, "3-d portal image analysis in clinical practice: An evaluation of 2-d and 3-d analysis techniques as applied to 30 prostate cancer patients", Int. J. Radiat. Oncol. Biol. Phys. **46**, 1281-1290 (2000).

5 J. C. Stroom, M. J. J. Olofsen-van Acht, S. Quint, M. Seven, M. De Hoog, C. L. Creutzberg, H. C. J. De Boer, and A. G. Visser, "On-line set-up corrections during radiotherapy of patients with gynecologic tumors", International Journal of Radiation Oncology Biology Physics **46**, 499-506 (2000).

6 J. Van de Steene, F. Van den Heuvel, A. Bel, D. Verellen, J. De Mey, M. Noppen, M. De Beukeleer, and G. Storme, "Electronic portal imaging with on-line correction of setup error in thoracic irradiation: Clinical evaluation", International Journal of Radiation Oncology Biology Physics **40**, 967-976 (1998).

7 G. Soete, J. Van de Steene, D. Verellen, V. Vinh-Hung, D. Van den Berge, D. Michielsen, F. Keuppens, P. De Roover, and G. Storme, "Initial clinical experience with infrared-reflecting skin markers in the positioning of patients treated by conformal radiotherapy for prostate cancer", Int. J. Radiat. Oncol. Biol. Phys. **52**, 694-698 (2002).

8 L. Court, I. Rosen, R. Mohan, and L. Dong, "Evaluation of mechanical precision and alignment uncertainties for an integrated ct/linac system", Med. Phys. **30**, 1198-1210 (2003).

9 J. Kress, S. Minohara, M. Endo, J. Debus, and T. Kanai, "Patient position verification using ct images", Med. Phys. **26**, 941-948 (1999).

10 S. Ryu, F. F. Yin, J. Rock, J. E. Zhu, A. Chu, E. Kagan, L. Rogers, M. Ajlouni, M. Rosenblum, and J. H. Kim, "Image-guided and intensity-modulated radiosurgery for patients with spinal metastasis", Cancer **97**, 2013-2018 (2003).

11 L. Pisani, D. Lockman, D. Jaffray, D. Yan, A. Martinez, and J. Wong, "Setup error in radiotherapy: On-line correction using electronic kilovoltage and megavoltage radiographs", Int. J. Radiat. Oncol. Biol. Phys. **47**, 825-839 (2000).

[12] M. J. Murphy, "An automatic six-degree-of-freedom image registration algorithm for image-guided frameless stereotaxic radiosurgery", Med. Phys. **24**, 857-866 (1997).

[13] L. Lemieux, R. Jagoe, D. R. Fish, N. D. Kitchen, and D. G. Thomas, "A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs", Med. Phys. **21**, 1749-1760 (1994).

[14] W. Birkfellner, J. Wirth, W. Burgstaller, B. Baumann, H. Staedele, B. Hammer, N. C. Gellrich, A. L. Jacob, P. Regazzoni, and P. Messmer, "A faster method for 3d/2d medical image registration--a simulation study", Phys. Med. Biol. **48**, 2665-2679 (2003).

[15] D. Fu, G. Kuduvalli, V. Mitrovic, W. Main, and T. Larry, "Automated skull tracking for the cyberknife (r) image-guided radiosurgery system", Medical Imaging 2005, Proc. of SPIE **5744**, 1605-7422 (2005).

[16] J. Pouliot, A. Bani-Hashemi, J. Chen, M. Svatos, F. Ghelmansarai, M. Mitschke, M. Aubin, P. Xia, O. Morin, K. Bucci, M. Roach, P. Hernandez, Z. R. Zheng, D. Hristov, and L. Verhey, "Low-dose megavoltage cone-beam ct for radiation therapy", International Journal of Radiation Oncology Biology Physics **61**, 552-560 (2005).

[17] D. A. Jaffray, "Emergent technologies for 3-dimensional image-guided radiation delivery", Sem. Rad. Onc. **15**, 208-216 (2005).

[18] D. A. Jaffray, J. H. Siewerdsen, J. W. Wong, and A. A. Martinez, "Flat-panel cone-beam computed tomography for image-guided radiation therapy", International Journal of Radiation Oncology Biology Physics **53**, 1337-1349 (2002).

[19] M. H. P. Smitsmans, J. de Bois, J. J. Sonke, A. Betgen, L. J. Zijp, D. A. Jaffray, J. V. Lebesque, and M. van Herk, "Automatic prostate localization on cone-beam ct scans for high precision image-guided radiotherapy", International Journal of Radiation Oncology Biology Physics **63**, 975-984 (2005).

[20] M. G. Herman, "Clinical use of electronic portal imaging", Semin Radiat Oncol **15**, 157-167 (2005).

[21] B. A. Groh, J. H. Siewerdsen, D. G. Drake, J. W. Wong, and D. A. Jaffray, "A performance comparison of flat-panel imager-based mv and kv cone-beam ct", Med. Phys. **29**, 967-975 (2002).

[22] R. Onimaru, H. Shirato, H. Aoyama, K. Kitakura, T. Seki, K. Hida, K. Fujita, K. Kagei, T. Nishioka, T. Kunieda, Y. Iwasaki, and K. Miyasaka, "Calculation of rotational setup error using the real-time tracking radiation therapy (rtrt) system and its application to the treatment of spinal schwannoma", Int. J. Radiat. Oncol. Biol. Phys. **54**, 939-947 (2002).

[23] K. Kitamura, H. Shirato, S. Shimizu, N. Shinohara, T. Harabayashi, T. Shimizu, Y. Kodama, H. Endo, R. Onimaru, S. Nishioka, H. Aoyama, K. Tsuchiya, and K. Miyasaka, "Registration accuracy and possible migration of internal fiducial

gold marker implanted in prostate and liver treated with real-time tumor-tracking radiation therapy (rtrt)", Radiother. Onc. **62**, 275-281 (2002).

[24] P. W. Chung, T. Haycocks, T. Brown, Z. Cambridge, V. Kelly, H. Alasti, D. A. Jaffray, and C. N. Catton, "On-line asi portal imaging of implanted fiducial markers for the reduction of interfraction error during conformal radiotherapy of prostate carcinoma", Int. J. Radiat. Oncol. Biol. Phys. **60**, 329-334 (2004).

[25] H. C. de Boer, J. R. van Sornsen de Koste, C. L. Creutzberg, A. G. Visser, P. C. Levendag, and B. J. Heijmen, "Electronic portal image assisted reduction of systematic set-up errors in head and neck irradiation", Radiother. Onc. **61**, 299-308 (2001).

[26] S. C. Erridge, Y. Seppenwoolde, S. H. Muller, M. van Herk, K. De Jaeger, J. S. Belderbos, L. J. Boersma, and J. V. Lebesque, "Portal imaging to assess set-up errors, tumor motion and tumor shrinkage during conformal radiotherapy of non-small cell lung cancer", Radiother. Onc. **66**, 75-85 (2003).

[27] A. L. Fielding, P. M. Evans, and C. H. Clark, "Verification of patient position and delivery of imrt by electronic portal imaging", Radiother. Onc. **73**, 339-347 (2004).

[28] S. Kumar, K. Burke, C. Nalder, P. Jarrett, C. Mubata, R. A'Hern, M. Humphreys, M. Bidmead, and M. Brada, "Treatment accuracy of fractionated stereotactic radiotherapy", Radiother. Onc. **74**, 53-59 (2005).

[29] K. W. Leszczynski, S. Loose, and S. Boyko, "An image registration scheme applied to verification of radiation therapy", Br. J. Radiol. **71**, 413-426 (1998).

[30] G. K. Matsopoulos, K. K. Delibasis, N. A. Mouravliansky, and K. S. Nikita, "Automatic medical image registration schemes using global optimization techniques", Stud Health Technol Inform **79**, 463-491 (2000).

[31] R. Bansal, L. H. Staib, Z. Chen, A. Rangarajan, J. Knisely, R. Nath, and J. S. Duncan, "Entropy-based dual-portal-to-3-dct registration incorporating pixel correlation", IEEE Trans. Med. Imag. **22**, 29-49 (2003).

[32] K. G. Gilhuijs, P. J. van de Ven, and M. van Herk, "Automatic three-dimensional inspection of patient setup in radiation therapy using portal images, simulator images, and computed tomography data", Med. Phys. **23**, 389-399 (1996).

[33] L. M. Sirois, D. H. Hristov, and B. G. Fallone, "Three-dimensional anatomy setup verification by correlation of orthogonal portal images and digitally reconstructed radiographs", Med. Phys. **26**, 2422-2428 (1999).

[34] B. Emami, J. Lyman, A. Brown, L. Coia, M. Goitein, J. E. Munzenrider, B. Shank, L. J. Solin, and M. Wesson, "Tolerance of normal tissue to therapeutic irradiation", International Journal of Radiation Oncology Biology Physics **21**, 109-122 (1991).

[35] J. T. Lyman, "Complication probability as assessed from dose volume histograms", Radiation Research **104**, S13-S19 (1985).

[36] C. P. Karger, D. Schulz-Ertner, B. H. Didinger, J. Debus, and O. Jakel, "Influence of setup errors on spinal cord dose and treatment plan quality for cervical spine tumours: A phantom study for photon imrt and heavy charged particle radiotherapy", Phys. Med. Biol. **48**, 3171-3189 (2003).

[37] B. C. J. Cho, M. van Herk, B. J. Mijnheer, and H. Bartelink, "The effect of set-up uncertainties, contour changes, and tissue inhomogeneities on target dose-volume histograms", Med. Phys. **29**, 2305-2318 (2002).

[38] W. Song, B. Schaly, G. Bauman, J. Battista, and J. Van Dyk, "Image-guided adaptive radiation therapy (igart): Radiobiological and dose escalation considerations for localized carcinoma of the prostate", Med. Phys. **32**, 2193-2203 (2005).

Chapter 2

Image Registration: Materials and Methods

2.1. Overview

The method for measuring patient position discussed in this thesis is based on the comparison between an orthogonal EPID image pair that represents the patient in treatment position and multiple orthogonal pairs of megavoltage DRR (MDRR) images, obtained from the CT data set used for treatment planning. Each MDRR image pair represents a certain patient orientation. The various components necessary to facilitate this procedure as well as the method used to validate its accuracy are outlined in detail in the subsequent sections of this chapter. The following paragraph provides an overview of the method.

Figure 2.1 shows a block diagram of the registration algorithm. The initial setup step consists of three parts:

- The data base of MDRR image pairs is computed. Each MDRR pair affords an AP and an LR view on the patient and represents a certain patient orientation (rotations in three dimensions about the treatment isocenter).

- Next, instructions on the search strategy are set, which determine the speed and accuracy of the method by specifying how many MDRR image pairs are to be compared with EPID image pairs. Information about which cost function is to be used, what image pre-processing steps are to be taken and where to locate files are also detailed.
- The EPID image pair is introduced into the algorithm. EPID images of the patient are taken in several treatment fractions: one image in anterior-posterior (AP) view and the other in latero-lateral (LR) view.

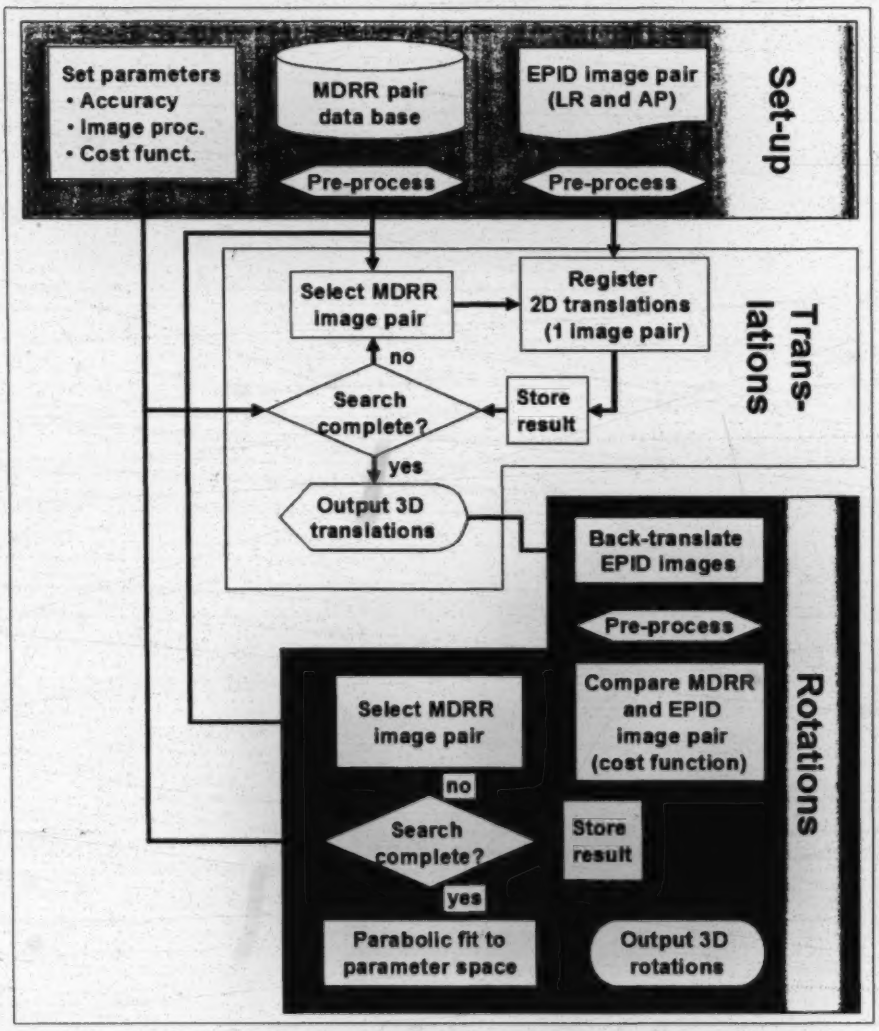

Figure 2.1 Flowchart outlining the 3D image registration method employed in this thesis

- The set-up step concludes with the pre-processing of both, the EPID and the MDRR images. The image processing steps employed as well as their significant influence on the registration accuracy are discussed in depth later in this thesis.

After the setup step, patient translations along each of three orthogonal axes are determined; the patient coordinate system is depicted in Figure 2.2. To this end, AP and LR view EPID images are registered separately with their corresponding MDRR images using a 2D registration routine (see chapter "2.7. Registering 3D Translations" for details).

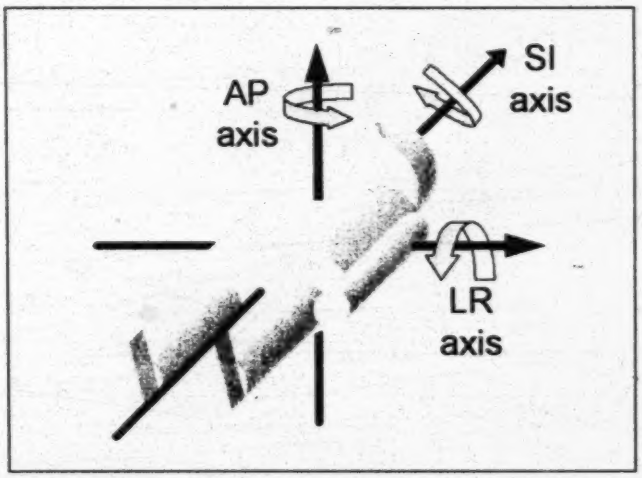

Figure 2.2 The patient coordinate system; translations along and rotations about each of the three orthogonal axes are registered with the method describe in this chapter.

The final step of the method provides registration of the rotations about the AP, SI (superior-inferior) and LR axes. This is accomplished by first aligning the beam axis in the EPID images with the planning isocenter (i.e. the beam axis in the MDRR images), which is done through back-translating the EPID images by the distances found in the previous (translations) registration step. Then the back-translated EPID image pair is compared with the MDRR image pairs in the data base. The degree of similarity is measured with a cost function, which assumes a minimal value for best correspondence. Since each MDRR image pair represents a certain 3D patient rotation, the MDRR image pair most similar to the EPID image pair is indicative of the patient's

rotation. Lastly, a parabolic fit around the minimum cost function value in the parameter space of rotation angles is applied. The location of the minimum of this parabola is the registration result for patient rotations.

2.2. Hard- and Software

All calculations, except for the generation of the MDRR data base and dose calculations, were carried out on a Windows XP workstation equipped with a Xeon processor (2 GHz, 533 MHz front side bus) and 1 GB random access memory (RAM).

The Pinnacle3 ® TPS (version 7.7a, Philips Medical Systems, Bothell, WA) running on a Sun Blade 1000 workstation (5 GB memory) was used to generate the MDRR data base and for dose calculations. Two sections of code that interfaced directly with the Pinnacle TPS ("plug-ins") were written in C++: one to automatically generate the MDRR data base ("MDRR-plug-in") and another to facilitate the computation and summation of dose distributions for multiple fractions of a patient's treatment ("dose-plug-in"). The registration framework code was written in Matlab, version 7.01 (© The Mathworks Inc. 2004). The framework performs the following tasks:

- Image pre-processing
- Searches of the MDRR data base
- Calling of the routine for registering 2D translations (via a DOS command)
- Evaluation of cost functions for registration of rotations
- Parabolic fit to the parameter space

Its various control parameters, such as options for the search depth of the MDRR data base and for the display of images, errors, warnings and progress information are specified in an initialization file that is read by the framework. Registration results are written into text files by the framework, using a custom format that is easily readable by humans as well as evaluation code. This code, also written in

Matlab was used to summarize and evaluate the results of multiple registration runs in an expedient manner.

The 2D-registration routine used to register translations was based on ansi-C code published by Thevenaz *et al.*[1] (http://bigwww.epfl.ch/algorithms.html), which was modified for the purpose of this work and tied into the registration framework (see section "2.7.2 2D Translations").

2.3. Generation of the MDRR data base

Both, the humanoid phantom (Pixy, Radiology Support Devices Inc., Long Beach, CA, see Figure 2.3) and patients were imaged on a PQ 5000 CT (Philips Medical Systems, Bothell, WA) scanner at 130 kV, using 3.0 mm slice thickness, pitch 1.0, a pixel size of 0.9375 mm and a 512x512 pixel matrix per slice. The CT data sets were imported into the Pinnacle3 ® workstation. The MDRR-plug-in automatically generated pairs of MDRRs as beam's eye views (BEV) in the isocenter plane for a 40 x 40 cm^2 field of view (FOV).

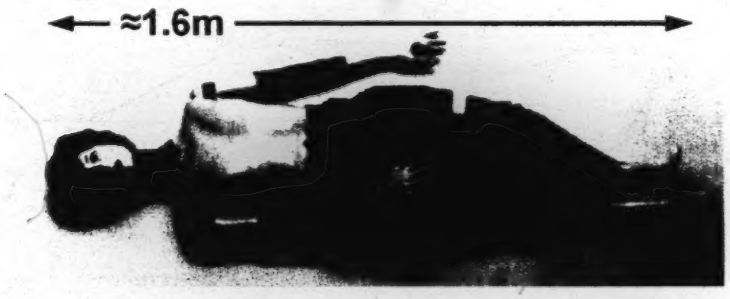

Figure 2.3 The humanoid phantom ("Pixy").

Digitally reconstructed radiographs (DRRs) are computed by numerically ray-tracing photon beams through a CT data set (Figure 2.5). Each volume element (voxel) within this three dimensional matrix contains photon attenuation coefficients, expressed as CT numbers in Hounsfield units (HU): [2]

$$CTnumber = 1000 \cdot \frac{\mu - \mu_{water}}{\mu_{water}}; \qquad (2.1)$$

where μ is the linear attenuation coefficient of the material in a given voxel during acquisition of the CT scan and μ_{water} the linear attenuation coefficient of water.[3] μ contains the electronic cross sections for coherent (Raleigh) scattering, photo effect, Compton scattering and pair production between a photon and an electron within the object being CT scanned: [3]

$$\mu = \frac{\rho \cdot N_A \cdot Z}{A}\left(\sigma_{photoeffect} + \sigma_{coherent} + \sigma_{compton} + \sigma_{pairproduction}\right) . \qquad (2.2)$$

Here, the factor of material density ρ, Avogadro constant N_A, proton number Z and atomic number A represents the electron density of the material. All of the four cross sections σ depend strongly on the photon energy E within the range of up to 2 MeV which is most frequently used for medical imaging; we therefore have $\mu = \mu(E)$. The CT numbers found by the CT scan at an average energy of approx. 60 keV are therefore different for the ones 'seen' by the linac beam when acquiring the EPID images at an average energy of approx. 2 MeV. This results in a marked difference in image appearance, as shown in Figure 2.4. To generate DRRs that look more like the megavoltage EPID images obtained with a linear accelerator (MDRRs), μ (which is derived from the CT number (Equation 2.1) obtained in the CT scan for each voxel) was scaled to the value it would have for this higher energy. Since this adjustment of μ depended on the unknown material present in each voxel the scaling process could only be an approximation. It is built into the Pinnacle TPS and could be adjusted to achieve best visual agreement between the MDRRs and an actual EPID image (Figure 2.4).

Figure 2.4 Pinnacle's scaling parameter for attenuation coefficients was set to reach best visual agreement between MDRR image (left and middle) and EPID image (right). Both MDRRs were generated from the same CT data set, left scaled to low (keV range) and right to high (MeV range) energy.

The algorithm performing the actual ray tracing is also part of the Pinnacle software. It is an approximation as well, because it does not attempt to accurately model photon interactions in the patient volume; specifically, scattering of photons within the patient volume, which increase noise and reduce contrast in the EPID images, is ignored in the resulting MDRRs.

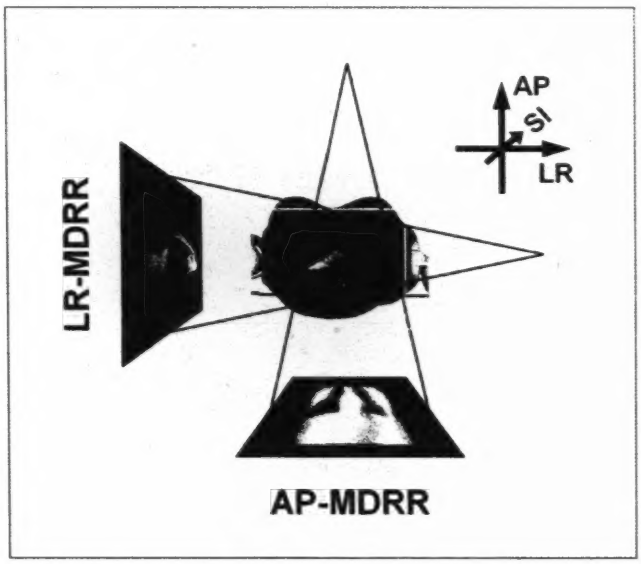

Figure 2.5 MDRR images are generated by ray-tracing through the CT data set, which is depicted in the center (here part of the torso of the humanoid phantom). AP images are calculated with simulated beams that originate from a source above the CT data set and are propagated trough the CT data set from top to bottom. LR images are the result of simulated beams traversing the CT data set from right to left.

MDRRs are created as orthogonal pairs (Figure 2.5) in AP and LR view. They are calculated for different patient orientations, i.e. 3D rotations of the patient about the isocenter. This rotation can be achieved by either rotating the CT data set or, alternatively, the beam source around the patient. In this work the second alternative was utilized because the Pinnacle TPS did not allow translation or rotation of the CT data set. The simulated rotations of the patient (in the remainder of this section, the word 'patient' denotes both, the phantom as well as a real patient) about the three principal axes were transformed into equivalent couch, gantry and collimator angles of a medical linear accelerator, which represent the orientation of the beam source. The

couch, gantry and collimator angles that afford the beams-eye-view (BEV) corresponding to pre-defined rotations of the patient about his AP, SI and LR axis were determined as described in Appendix A.1.

Each pair of MDRR images thus represents a patient rotation vector $\Phi = (\Phi_{AP}, \Phi_{SI}, \Phi_{LR})$ containing rotation angles about the *AP*, *SI* and *LR* axis respectively ("pitch-roll-yaw" convention). This angle information was stored in a lookup table. The MDRRs covered a parameter space of angles ranging from -5° to +5° with an angular spacing of 0.5° for a total of 21 angles per axis. The total number of image pairs (i.e. the size of the parameter space) in the MDRR data base was therefore 9261. Each 40 x 40 cm^2 image consisted of 256 x 256 pixels (pixel size of 1.56 mm) and each pixel was represented by two bytes. Therefore the size of one patient's MDRR data base was 2.26 GB; it took 14 hrs to generate on the Sun Blade 1000 workstation.

2.4. Acquisition of the EPID Image Pairs

EPID images of patients and the humanoid phantom in treatment position were acquired with an aS 500 amorphous silicon flat panel imager, installed on a 2100EX linear accelerator (Varian Oncology Systems, Palo-Alto, CA). In the EPID, [4,5] incoming high energy photons interact predominantly through Compton Effect with a buildup layer consisting mostly of a copper sheet of 1 mm thickness, in which charged particles are generated. These then deposit dose in an underlying phosphor layer (340 μm of gadolinium oxysulfide, Gd_2O_2S:Tb), causing the production of optical photons at 540 nm (emission peak). These photons are detected by the light-sensitive pixel matrix which is layered underneath the phosphor and consists of amorphous silicon (a-Si). An alternative technology is the use of a charge coupled device (CCD) as light sensitive element. However, the small dimension of the CCD chip (less than approx. 5x5 cm^2) requires optical reduction of the image size, a process that suffers from low efficiency and results in a more bulky device. The EPID's a-Si matrix consists of 512 x 384 pixels with a physical size of 0.784 mm. Images were exposed at 6 MV with a field size of 20 x 20 cm^2. Images of the Pixie phantom were acquired with 1 monitor unit (MU) per image. Patient images were acquired with 1 or 2 monitor units. The source-to-object

distance was 100 cm and the source-to image distance was 145 cm. Two patient images were scaled slightly differently, probably because the image was taken while the EPID imager was not in its fully extended position. Based on the graticule pattern visible on the images, they were rescaled by a magnification factor of 0.97 and 0.95, respectively.

An orthogonal image pair of patients consisted of one AP view at 0° gantry angle and an LR view at either 90° or 270° gantry angle (Figure 2.6). The images of the Pixy phantom were also orthogonal pairs. Their purpose was to test the accuracy of the registration method. In order to test the method's ability to find patient rotations, EPID images were obtained for different rotations of the phantom around its AP, SI and LR axes. However, because of the flexibility and large size and weight of the phantom, rotating it in a precise manner would be impractical. Phantom rotations were therefore emulated by setting couch, gantry and collimator angles in such a way that the BEV of the phantom was the same as if the phantom itself had been rotated. For example -2° roll of the patient about the SI axis (Φ = (0, -2, 0)) would require EPID images at 2° (AP view) and 92° (LR view) gantry angle (Figure 2.7).

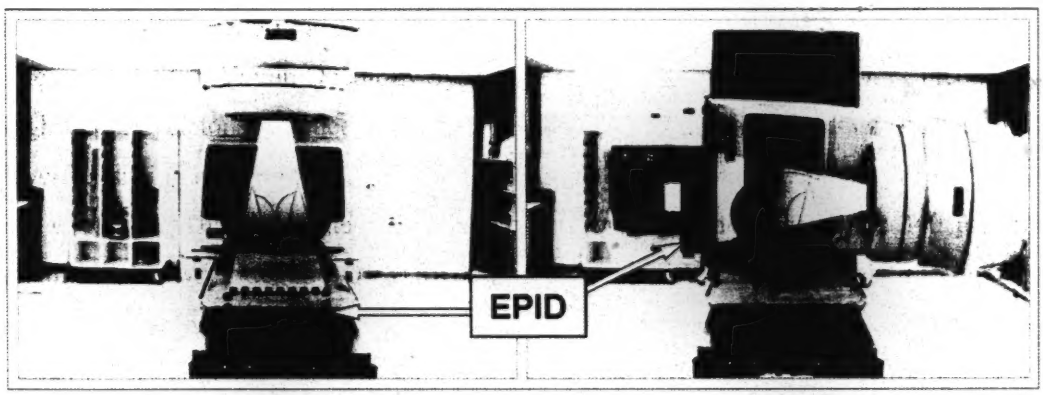

Figure 2.6 Acquisition of orthogonal EPID images of patients; AP view (left) and LR view (right).

For the general, three-dimensional case the couch, gantry and collimator angles corresponding to those known phantom rotations were calculated as outlined in Appendix A. This approach permits acquisition of EPID images of 'simulated' phantom rotations with the accuracy of the linac's rotation (0.1° deg for gantry and collimator, 0.2° for couch). It neglects the fact that the treatment couch is somewhat radio opaque:

whereas in the case of an actual phantom rotation the relative orientation of phantom and treatment couch would change, this is not the case here. In AP view, this will cause the couch support structure ('tennis racket') under the phantom to appear un-rotated relative to the phantom. However, this structure was only faintly visible on the images. Any influence it had on the registration would serve to degrade the results, so that the results reported represent a worse case scenario.

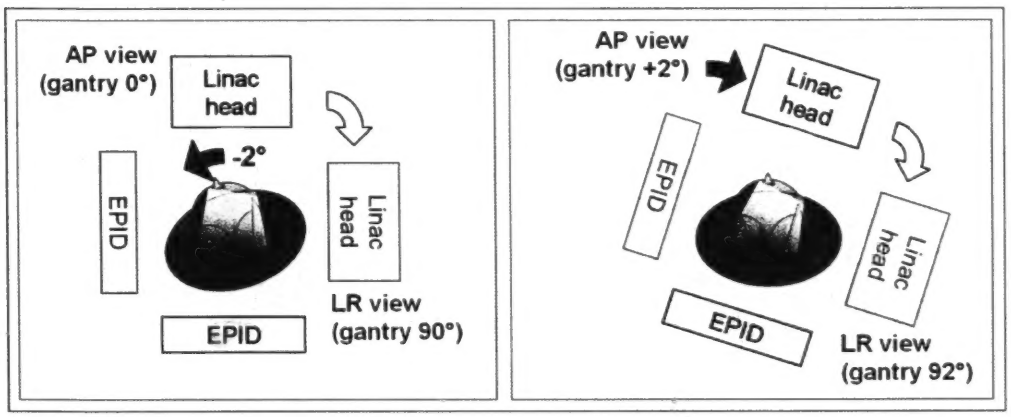

Figure 2.7 Equivalency of phantom roll around SI axis rotation (left) to the adjusted gantry angle (right) that is used to emulate rotations of the humanoid pixie phantom. The phantom rotations about the other axes wre similarly realized using combinations of couch, gantry and collimator angles of the linear accelerator.

2.5. Image pre-processing

Both, EPID and MDRR images underwent a series of pre-processing steps prior to registration. Both image types were re-sampled to the same physical pixel size and cropped to the desired FOV for registration (see Figure 2.8). At this point, further manual segmentation of image features could have take place. However, it was purposely avoided in order to render the results as independent as possible from individual users' judgments and maintain a high degree of automation of the method. Image histograms were aligned by means of shifting and stretching or compressing such that all pixel values fell within a pre-defined interval.[6] In this work, this histogram alignment step is preceded by two less common image processing steps: unsharp

masking and histogram matching.[7,8] Care was taken to make these steps universal in the sense that they need not be manually adapted to a changing imaging situation. For easier reference, the combined three steps of unsharp masking, histogram matching and histogram alignment are designated as *extended* image pre-processing, as opposed to *basic* image pre-processing, here defined as only histogram alignment. The image pre-processing steps are illustrated in Figure 2.9 and outlined below.

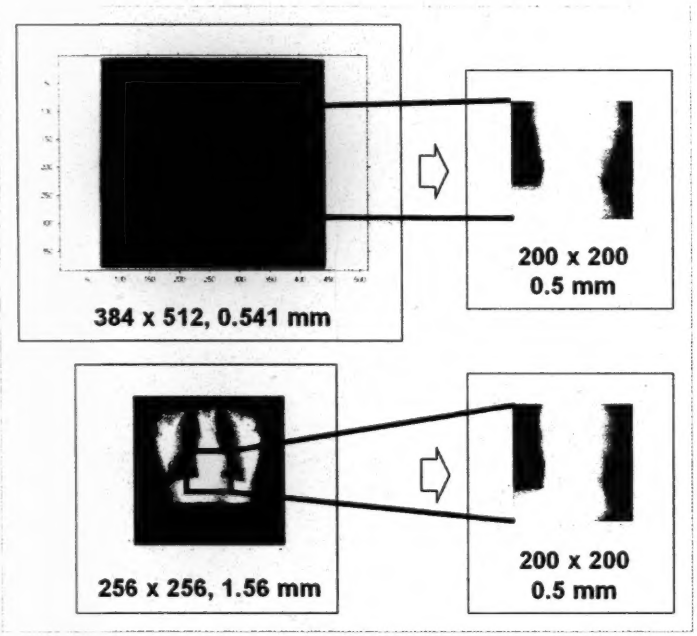

Figure 2.8 The EPID image (top) and the MDRR image (bottom) are cropped and re-sampled to the same pixel size; shown here is only the AP view.

- Unsharp masking: A strongly blurred version of an image was subtracted from the original image. Then the mean value of the original image was added. The blurring was performed with a fairly wide Gaussian kernel: the kernel width was one-half of the width of the image and the Gaussian width (2 σ) was one-fourth of the image width. This step serves as a high pass and reduces the magnitude of image features at low spatial frequencies in the image.

- Histogram matching: To enhance image contrast, each image's histogram was then matched to a normal distribution of width $\sigma = 0.2$ times the dynamic range of the original image (i.e. the frequency of occurrences of pixel values was matched to that of the normal distribution).
- Histogram alignment: Histograms were re-scaled (stretched or compressed) and then shifted to lie in the arbitrary but constant interval of [0 255]. The data format of the image pixel values remained 'float'.

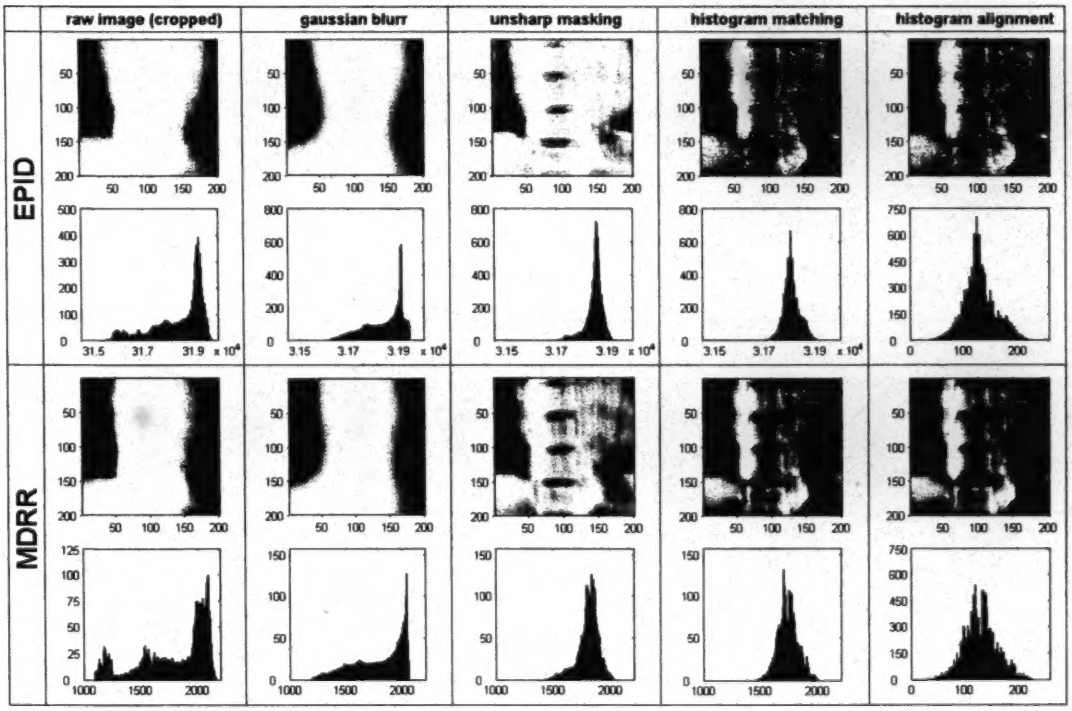

Figure 2.9 Image pre-processing steps used in the method. The original (raw), cropped EPID and MDRR images (first column); Original images convolved with a large Gaussian kernel (second column); Blurred images subtracted from original images (third column); Histograms matched to a normal distribution (fourth column); Histograms aligned to same interval (fifth column). After unsharp masking overall image intensities are more balanced and finer structures (e.g. vertebral bodies) become visible. Histogram matching enhances contrast and further enhances similarity. Finally the image histograms are shifted into the same gray-value range of [0, 255]; pixel values are non-integer (double). For consistent visualization, all images shown have been scaled to gray values between 0 and 255.

In addition, the EPID images were corrected for:

- The magnification resulting from the ratio of the source-to-imager distance (SID = 145 cm) to the source-to-axis distance (SAD = 100 cm).
- The position of the beam axis on the EPID imager: Since the beam axis was not exactly centered on the imager, images needed to be shifted to be centered on the beam axis. This procedure was performed differently for patient and phantom images:
  i. Patient images are acquired by placing a graticule filter in the beam, which generates a cross pattern of bright dots to visually judge a patient's position (Figure 2.10). The beam axis coincides with the central dot of the cross pattern. The center of that dot was found by applying a Canny edge detection filter to the region around the dot. Once the pixels constituting the edge of the central dot were determined, the center-of-mass (CM) of those pixels was calculated. The coordinates of the CM are the position of the beam axis on the EPID imager.

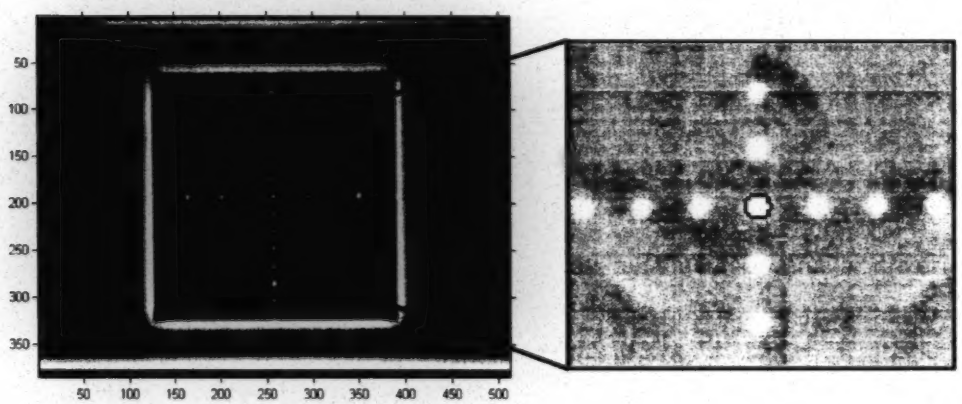

Figure 2.10 The position of the beam axis on the EPID image is determined by the cross shaped pattern generated by the graticule filter (left). It is found by calculating the center of mass (CM) coordinates of the edge pixels of the central dot, visible in the magnified image on the right.

ii. If no graticule filter was used (while testing the algorithm using the humanoid phantom), the position of the beam axis on the portal imager was determined as the CM of the edges of a 10x10 cm^2 open field image as displayed in Figure 2.11. Because of sagging of the portal imagers at 90° or 270° gantry angle (by approximately 2 mm), this correction was done separately for the AP and LR view images.

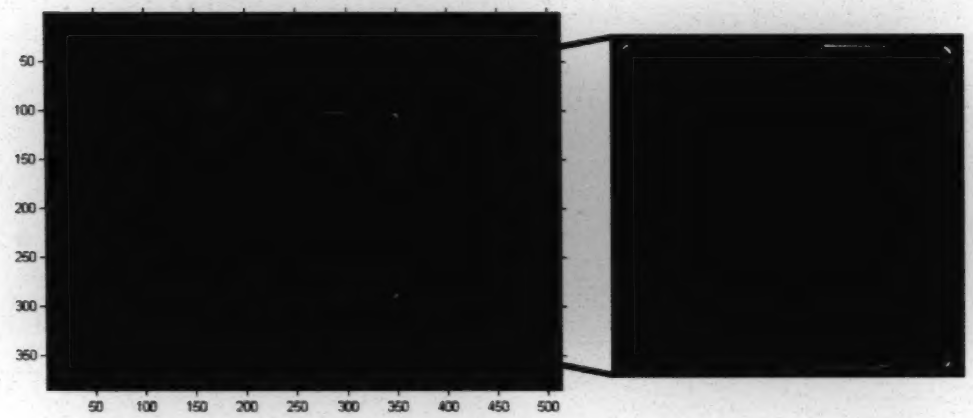

Figure 2.11 Epid image of a 10 x 10 cm^2 field (left). The magnified image on the right shows the edge pixels of the square field (black). The beam axis on the EPID image was found by calculating the CM of those edge pixels. The images here are displayed in false color for increased visual contrast.

- Rotational offset of the imager: It was further observed that the straight collimator edges of a 10 x 10 cm^2 open field image did not run perfectly parallel to the EPID's pixel matrix. These 'breaks' in the edge pixels of the image, visible in Figure 2.12, were due to a slight rotation of the EPID imager relative to the beam-defining collimator jaws. This rotational offset of the imager was 0.25° at 0° gantry angle and 0.35° deg at 90° gantry angle. It was accounted for by rotating the corresponding EPID images by the same amount prior to further analysis.

- Graticule filter pattern: Clinical patient images are generated by placing a graticule filter in the beam, which generates a cross pattern of bright

dots in the EPID images to visually judge a patient's position. This pattern is not part of the corresponding MDRR images. To minimize the negative influence of the graticule pattern on the registration process, this pattern was – after measuring the beam axis position – largely removed by replacing the dots with the average value of their surrounding neighborhood (see Figure 2.13).

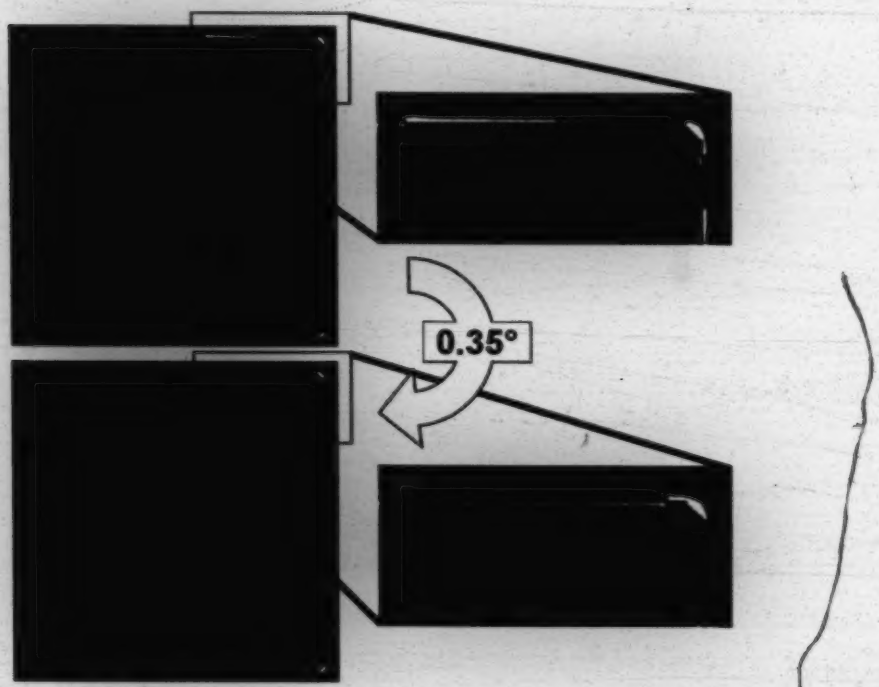

Figure 2.12 Rotational offset of the EPID images was found by exposing a square field. Initially, the edge pixels (black) do not form a straight line (top images, marked by arrows). After rotating the image by 0.35°, the edge pixels form a straight line. The images here are displayed in false color for increased visual contrast.

The image pre-processing steps took less than two seconds per image. However, since up to 1000 MDRR image pairs needed to be pre-processed in order to register one pair of orthogonal EPID images (while searching the MDRR data base using the smallest angular interval between consecutive MDRRs as described in section "2.7. Registering 3D Translations"), this step could lengthen the registration procedure

by more than an hour. Pre-processed images were therefore stored in a temporary folder that made repeat-processing of the same image unnecessary.

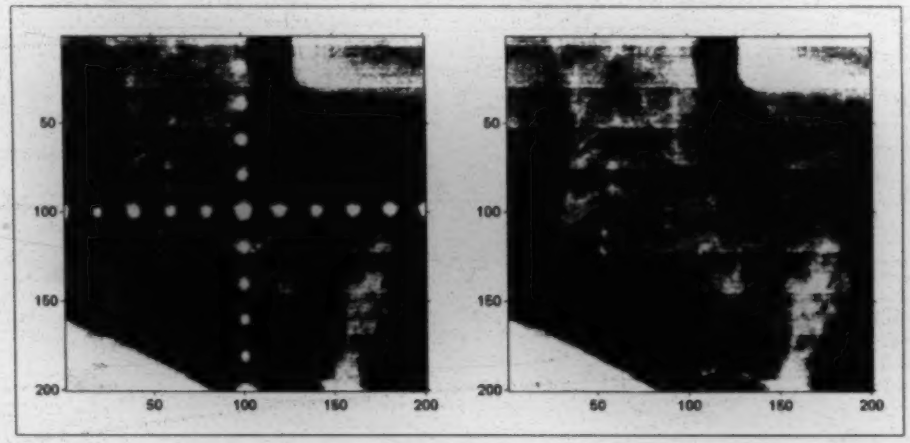

Figure 2.13 Part of the pre-processing of clinical EPID images is a routine that automatically removes the cross pattern of white dots (left image), which are caused by the gantry-mounted graticule filter. The routine replaces pixel values within the dots with local averages around them (right image).

2.6. Cost Functions

Cost functions describe the degree of similarity of two images. Three different cost functions that are widely used in medical image registration were tested in this work for their ability to provide accurate registration results. They are based on well known pixel value similarity measures and scaled to exhibit a minimum for best image alignment. They were evaluated for the overlapping pixels between the two images being registered. In the following, let I and J be two images of identical number of pixels M x N, indexed by m and n so that $I_{m,n}$ and $J_{m,n}$ are the pixel values at position (m, n) in images I and J respectively. $\bar{I}$ and $\bar{J}$ denote the average pixel value of each image ($\bar{I} = \frac{1}{M \cdot N} \sum_{m,n} I_{m,n}$). The cost functions, CF, are then defined as follows, based on:

- Mean squared difference of pixel values:

$$CF_{MS} = \frac{1}{M \cdot N} \sum_{m,n} \left(I_{m,n} - J_{m,n} \right)^2 . \quad (2.3)$$

It can be shown that this cost function works best when the two images being compared only differ by Gaussian noise[9]. Because of its square weighting of the pixel values, CF_{MS} is very sensitive to a small number of pixels that differ by a large amount in images I and J. CF_{MS} is the similarity measure native to the Marquardt-Levenberg algorithm used to register 2D translations (section "2.7. Registering 3D Translations").

- Correlation-coefficient (also known as Pearson's linear correlation coefficient):[8]

$$CF_{CC} = 1 - \frac{\sum_{m,n} \left[\left(I_{m,n} - \bar{I} \right) \left(J_{m,n} - \bar{J} \right) \right]}{\sqrt{\sum_{m,n} \left(I_{m,n} - \bar{I} \right)^2 \sum_{m,n} \left(J_{m,n} - \bar{J} \right)^2}} . \quad (2.4)$$

This cost function is invariant under the transformation $I \rightarrow a \cdot I + b$, i.e. when scaling the images intensities linearly and adding a constant offset: $CF_{CC}(I,J) = CF_{CC}(aI + b, cJ + d)$. It is therefore more robust than CF_{MS}.

- Normalized mutual information:[10]

$$CF_{MI} = 1 - \frac{2 \cdot \left[H(I) + H(J) - H(I,J) \right]}{H(I) + H(J)} . \quad (2.5)$$

Here, $H(I)$ and $H(J)$ are the Shannon entropies[11] of images I and J, respectively: $H(I) = -\sum_i p_i(I) \cdot \log(p_i(I))$. $p(I)$ is the probability distribution of gray values of image I and the index i denotes a histogram bin of image I. $H(I,J)$ is the joint entropy of images I and J. It is defined similar to $H(I)$, but $p(I)$ is replaced by $p(I,J)$, the histogram of joint probabilities of gray values of images I and J: $H(I,J) = -\sum_j \sum_i p_{i,j}(I,J) \cdot \log(p_{i,j}(I,J))$. The entropy H of an image is a measure for its information content. If, for example, the

image were to consist of a uniform gray (all pixels have the same value and thus the image contains no information), $p_i = 1$ for one index i (i.e. the index labeling the histogram bin containing the gray value that all the image pixels possess) and $p_i = 0$ all other i and therefore $H(I) = 0$. If pixel values were randomly distributed, ($p_i(I) = 1/MN$ for all i), the entropy would be $H(I) = -MN \cdot (1/MN \cdot \log(1/MN)) = \log(MN)$. The joint entropy $H(I,J)$ measures how far one of the images explains the other. If both images are identical, the second image explains the other perfectly and adds no more information to the entirety of both images. The joint histogram of probabilities of images I and J is then equal to the histogram of their individual probabilities: $p(I,J) = p(I) = p(J)$. Only the terms with $i = j$ contribute to the double sum and the joint entropy $H(I,J) = H(I) = H(J)$ takes on a minimal value. If, on the other hand, both images are wholly unrelated, the joint entropy equals the sum of the entropies of both individual images. In general, $H(I,J)$, $H(I)$ and $H(J)$ are related by the inequality: $\max(H(I),H(J)) \leq H(I,J) \leq H(I) + H(J)$. The values for CF_{MI} as defined above therefore vary between 0 for identical images and 1 for totally unrelated images I and J.

2.7. Registering 3D Translations

2.7.1 Searching the MDRR data base

Of all the MDRR image pairs in the data base, the one that corresponds most closely with the EPID image pair (lowest cost function) will provide the truest value for translations. A simple interval shrinking scheme was therefore used to search the MDRR data base for this pair: first, a given EPID image pair was 2D-registered with all the MDRR image pairs that differed by an angle interval of 2.5° and the value of the cost function achieved in each case was stored. Thus, the full angular range from -5° to +5° was covered for each axis. The angles corresponding to the MDRR pair that

achieved the lowest cost function (the best match) now served as the starting point for the next iteration. This next iteration consisted of selecting MDRRs that are separated by an angle interval of 1° within an angular range of ±2° around the previous minimum. Those were again 2D-registered and the process was repeated with the finest interval (0.5°) for an angle range from -1° to +1°. The translations corresponding to the smallest cost function value were reported as final result. This strategy reduced the number of necessary 2D-registrations from 18522 which would have been necessary to cover the full parameter space, to approximately 1200.

2.7.2 2D-translations

The AP and LR view EPID images were registered separately with their corresponding MDRR images to determine their 2D translations. As mentioned above (chapter "2.2. Hard- and Software"), 2D registration between one EPID and one MDRR image was performed with code written in ansi-C by Thevenaz *et al.*[1]. The code attempts to transform a test image J into a reference image I by applying a transformation $T(J)$. In our case, the code was capable of performing a 3D linear conformal transformation (rotations, translations and scaling). Since our registration task involves only two-dimensional images and at this stage no rotations, the code was restricted to find only translation in two dimensions. The transformation T is therefore parameterized only by the distances t_x and t_y of those translations: $T(J) = T(J;t_x,t_y)$. The code then attempts to minimize the cost function $CF = CF(I,T(J;t_x,t_y))$ between image I and the transformed image J by varying the values for t_x and t_y. This search through the space of possible values for t_x and t_y is done iteratively by the Levenberg-Marquardt (LM)[12] method, which is improved gradient search strategy. Applying the transformation T to image J involves interpolations of pixel values if translations other than multiples of the pixel size are to be considered. The code uses a cubic (spline) interpolation for this purpose. Lastly, to increase registration speed and robustness, registration is performed in a coarse-to-fine manner: the images are initially re-sampled at 1/8 of their native resolution and registered; the registration result is then used as start value for registration at 1/4 of the native resolution; this process is repeated at 1/2 and finally at the full native resolution.

The cost function used natively by the LM method is CF_{MS}. The code was expanded by adding routines for the calculation of CF_{CC} and CF_{MI}. Finally, the code was compiled, linked and the executable tied to our registration framework. Various initialization parameters were read from a custom-formatted text file and run-time parameters (image names, locations, sizes and cost functions) were passed to the executable via the DOS command calling it from the Matlab framework. The registration results (translations t_x, t_y as well as the value of the cost function) were passed back into the Matlab framework, again via the DOS interface. Registration of one image pair of 200 x 200 pixels takes approximately one second.

2.7.3 3D-translations

The value for translation along the LR axis was obtained from the 2D registration of the AP view MDRR and EPID images. Translations along the AP axis resulted from the LR view registration. Translations along the SI axis were obtained from both, the LR and AP view registration. The mean of those two values for the SI translation, which deviated by less than 0.25 mm (1SD) from each other, is reported by the routine (see Figure 2.14).

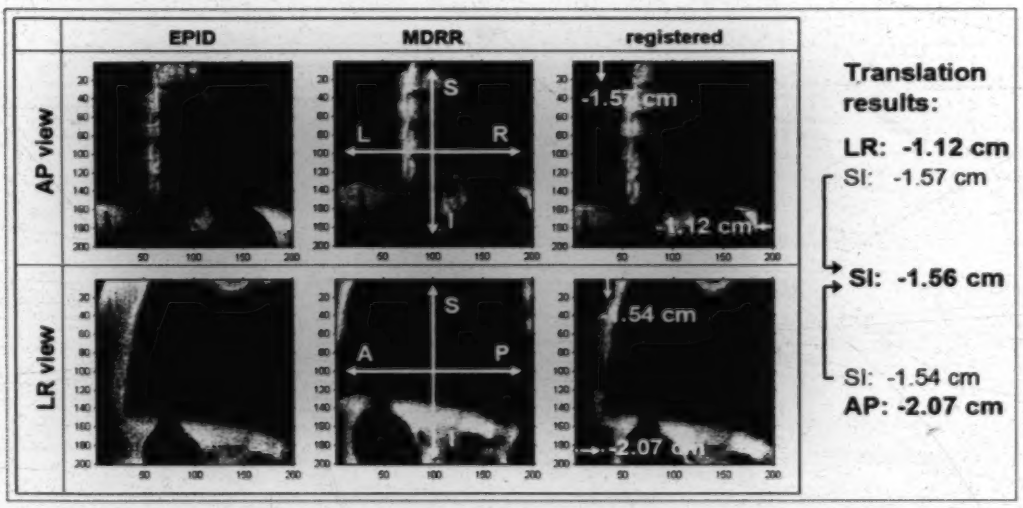

Figure 2.14 The 2D registration is performed separately for LR and AP view images. The AP view registers translations in S-I and L-R direction, whereas the LR view is sensitive to translations in S-I and A-P directions. The S-I translation is the overage of the two results obtained from each of the two views.

2.8. Registering 3D Rotations

Since each MDRR corresponds to a certain patient orientation (rotation vector Φ), the best matching MDRRs determined above (section "2.7. Registering 3D Translations") provide both the patient's translations and rotation angles Φ. The average of the MDRR angles for the (separately registered) AP and LR view were reported by the algorithm described by Sirois *et al.*[13]; their strategy for finding the 3D rotations is designated here as "*From translations*" for later reference. In this thesis, additional steps purely dedicated to determine more accurate rotation angles are introduced: first the raw EPID images are shifted by the translation values found above, thus aligning the beam isocenter in the EPID image with that in the planning CT data set. Then the raw EPID images are again pre-processed and finally the AP and LR view EPID images are simultaneously registered with the MDRR image pairs in the data base. This registration involves simply the computation of the cost function between the EPID image pair and the MDRR image pairs, again employing the interval shrinking strategy to search the MDRR data base as described in section "2.7. Registering 3D Translations". The patient's rotations could now be determined by simply reporting the rotations Φ corresponding to the best-matching (smallest cost function value) MDRR image pair; this strategy will be referred to as "*MDRR angle*".

Since any extremum can be approximated by a Taylor expansion to the second (quadratic) order, it was hypothesized that fitting a thee-dimensional parabola to the values of the cost function in an interval around this minimum would provide more accurate registration results. The second order Taylor expansion in three dimensions has the general form:

$$CF = a + \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}^T \cdot \begin{pmatrix} \phi_{AP} \\ \phi_{SI} \\ \phi_{LR} \end{pmatrix} + \frac{1}{2} \cdot \begin{pmatrix} \phi_{AP} \\ \phi_{SI} \\ \phi_{LR} \end{pmatrix}^T \cdot \begin{pmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{pmatrix} \cdot \begin{pmatrix} \phi_{AP} \\ \phi_{SI} \\ \phi_{LR} \end{pmatrix} ; \qquad (2.6)$$

$$= a + \vec{b}^T \vec{\phi} + \frac{1}{2} \vec{\phi}^T C \vec{\phi}$$

where T denotes the transpose operation and the constants a, $\vec{b}$ and C are determined

by a numerical least square fit to the measured cost function values around the minimum. The matrix C contains the second derivatives of our Taylor expansion; it is the Hessian Matrix of the parabola. Setting the first derivative of Equation 2.6 equal to zero ($0 = \vec{b}^T + C\vec{\phi}_{\min}$) yields the parabola's minimum. It is calculated by numerically determining the inverse C^{-1} of C, so that

$$\vec{\phi}_{\min} = -C^{-1} \cdot \vec{b} \quad . \tag{2.7}$$

To interpolate the parameter space with a parabolic fit, first the values of the cost function for MDRR pairs representing angles around that minimum (found by the "*MDRR angle*" strategy) were calculated with the help of a cost function. The MDRR angles usually lay in an interval of ±1° around the minimum in each axis of rotation, but larger intervals were used for coarser sampling of the parameter space. The parabola was then fit to those cost function values (Figure 2.15). The coordinates of the minimum of this parabola are the 3D rotation values reported by the method. This last strategy will be called "*Parabolic fit*" when referred to in the next chapter were results will be presented.

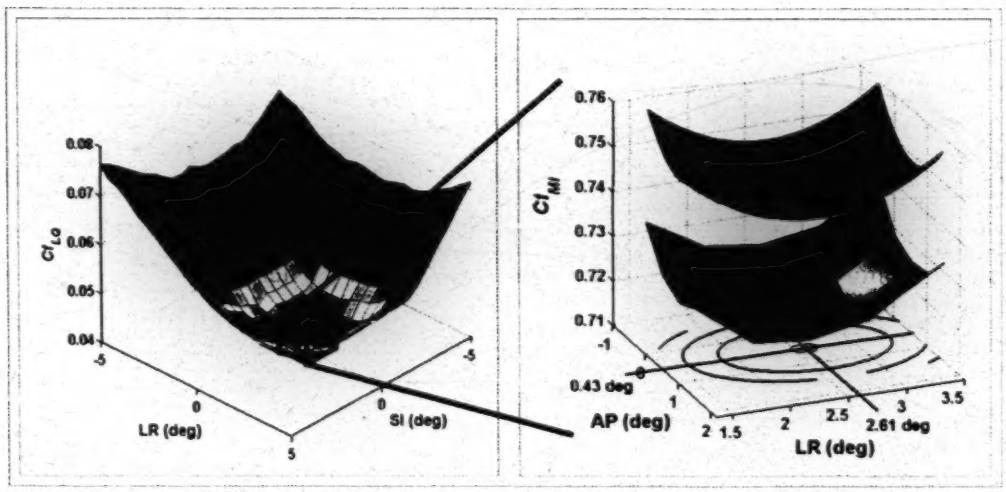

Figure 2.15 Two-dimensional cross section of the angle parameter space. Left the full parameter space, right the values around the minimum, to which a parabola is fitted. For display purposes, the parabola has been elevated along the vertical axis above the surface of the cost function.

As with any registration method, the possibility that the algorithm does not converge has to be considered. The following criteria were used to reject a registration result:

- One or more of the rotation angles found lies outside the pre-defined angle interval (±5° for the purposes of this thesis)
- One or more eigenvalues of the Hessian matrix of the parabola, the matrix C in Equation 2.6, are negative, which would indicate a saddle-shaped extremum or a maximum.

2.9. Coupling of Rotations and Translations

The method discussed in this work registers translation and rotation independently of each other. However this independence strictly exists only for non-divergent (parallel) beams. When using the diverging beams of a linear accelerator, rotations and translations are coupled translations because projection images of a 3-dimensional object change not only when the object is rotated, but also when the object is translated relative to the beam source. First, the coupling of rotations to translations is discussed and the reverse case (coupling of translations to rotations) will be discussed at the end of this section. For the dependence of registration results for rotations on the translations present in the phantom or patient, the following two cases can be distinguished, both of which have been tested in this thesis.

- Translations parallel to the beam axis (orthogonal to the imaging plane): This type of translation will cause magnification changes in the imaging plane, but details in the object will not change their relative position. A minimal, if any measurable, dependence of the registration results for rotations on the amount of translation is expected.
- Translations orthogonal to the beam axis (parallel to the imaging plane): In this case, details at different depths within the object, when they are projected onto the imaging plane (EPID), will change their relative positions. Some dependence of registration results for rotations

on the amount translation is expected. A projection image of a 3-dimensional object that has been translated by a small amount can appear similar to a projection image obtained by a small rotation of the object about an axis orthogonal to the direction of translation. The registration algorithm then 'perceives' a pseudo rotation when no actual rotation has occurred, leading to incorrect registration results.

The magnitude of the coupling of rotations to translation described in the second bullet above is calculated as follows: Figure 2.16 illustrates how two details A and B, located at different depths within a square object (separated by a vertical distance *h*) change their relative position in the imaging plane when the object is either translated or rotated.

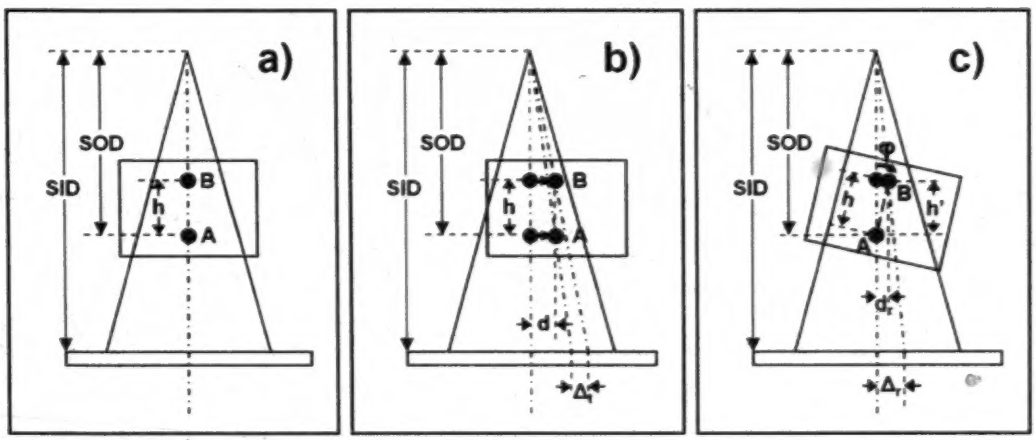

Figure 2.16 Geometry illustrating how the translation of an object (b) can appear similar to a rotation (c). Originally, both details A and B within the square object are aligned such that they over lap in the imaging plane. The details are separated in the case of both, translation (b) and rotation (c). SID = Source-to-Image-Distance, SOD = Source-to-Object-Distance

Originally, both details are positioned on the beam axis so that their projection images overlap in the imaging plane (Figure 2.16a). After the object is translated by a distance *d*, detail A moves a distance $d \cdot SID/SOD$ in the imaging plane and object B moves the distance $d \cdot SID/(SOD - h)$ (Figure 2.16b). The projections of both details in the imaging plane are now separated by the distance

$$\Delta_t = d \cdot \left(\frac{SID}{SOD - h} - \frac{SID}{SOD} \right). \tag{2.8}$$

If the object is rotated around detail A by an angle φ (Figure 2.16c), detail B also separates from detail A in the imaging plane, this time by a distance:

$$\Delta_r = d_r \frac{SID}{SOD - h'} = h \cdot \sin\varphi \cdot \frac{SID}{SOD - h \cdot \cos\varphi} \approx h \cdot \sin\varphi \cdot \frac{SID}{SOD - h}. \tag{2.9}$$

Equating Δ_r and Δ_t yields a relationship between the (apparent) rotation angle φ that results from a translation by a distance d:

$$\begin{aligned} \Delta_r = \Delta_t \quad &\Leftrightarrow \quad h \cdot \sin\varphi \cdot \frac{SID}{SOD - h} = d \cdot \left(\frac{SID}{SOD - h} - \frac{SID}{SOD} \right) \\ &\Leftrightarrow \quad \sin\varphi = \frac{d}{h} \cdot \frac{\left(\frac{SID}{SOD - h} - \frac{SID}{SOD} \right)}{\frac{SID}{SOD - h}} = \frac{d}{h} \cdot \left(1 - \frac{SOD - h}{SOD} \right) = \frac{d}{SOD} \end{aligned} \quad ; \tag{2.10}$$

so that:

$$\varphi = \sin^{-1}\left(\frac{d}{SOD} \right). \tag{2.11}$$

This result is independent of the distance h between the two details A and B and of the source-to-image distance SID. Figure 2.17 shows the apparent rotation angle as a function of displacement (translation) d for various values of SOD.

We also have to consider that the object is viewed in two orthogonal directions (AP and LR) with our method. While a translation along one axis, say, the LR axis, appears as rotation for the AP view, the LR view image is hardly changed, except for a very small change in magnification. If the image was registered solely on the basis of the AP view EPID image, the AP MDRR image best matching the translated image would be the one representing a rotation by an angle φ given by Equation 2.11. When we now also consider the LR view MDRR image corresponding to the angle φ (which is linked to the AP view image in our MDRR data base), we would find that it doesn't match well with the almost unchanged (apart from very small magnification changes)

LR view EPID image. The lowest cost function value will therefore be achieved with an MDRR image *pair* that represents a rotation angle between zero and φ as a compromise between the apparently rotated (AP view) and un-rotated (LR view) image pairs. To account for this effect, we include a factor 0.5 in Equation 2.11:

$$\varphi = 0.5 \cdot \sin^{-1}\left(\frac{d}{SOD}\right). \tag{2.12}$$

We will compare our experimental result with this equation.

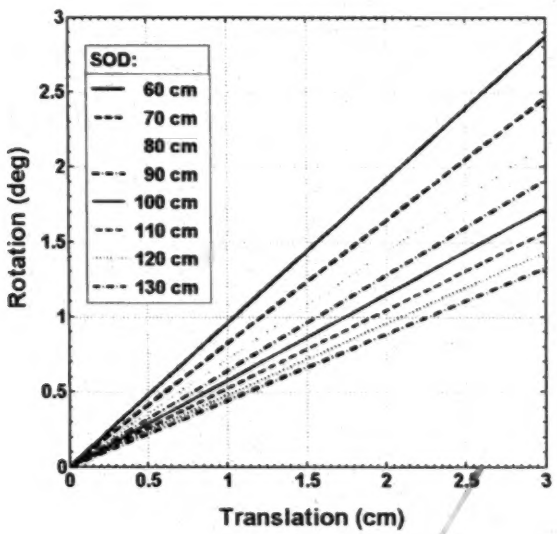

Figure 2. 17 Pseudo rotation angles as a function of the translation of an object, plotted for different SOD values according to Equation 2.11.

As can be seen in figure 2.16c, the rotation axis is orthogonal to both, the direction of translation and the beam (projection) axis. The registration method utilizes two orthogonal views onto the patient (AP and LR view), in which translations are visible. Table 2.1 summarizes which view 'sees' what translation along any of the three principal axes. It also lists the axis around which a pseudo rotation occurs as a result of the translation.

After the investigation of the dependence of rotations on translations, the reverse question could be asked: how far do registration results for translations depend on rotations? An orthogonal EPID image pair acquired of a rotated but not translated

phantom (or patient) shall be considered. The first step of our 3D registration consists of registering the 3D translations (see section "2.7. Registering 3D Translations"). In this step, the EPID image pair is compared with a number of MDRR image pairs, each representing a given rotation. The MDRR image pair whose rotation best matches the one captured in the EPID images will be the one yielding the lowest cost function value. This pair will therefore be used for the registration of the 3D translations, but it will not introduce an offset to the translational registration result because all MDRR images are generated without translations relative to the isocenter (they solely represent rotations about the isocenter). The strategy "*From translations*" (see section "2.8. Registering 3D Rotations") makes use of this fact and uses it to determine 3D rotations. We therefore expect not to detect an influence of rotations on translations. The dependence of translations on rotations will also be tested in this thesis by registering cases of a rotated by not translated phantom.

Table 2.1 Translations along any of the three orthogonal axes appear as rotations about axes orthogonal to the beam (projection) axis and to the direction of translation. Because orthogonal images are acquired (one in AP and one in LR view), translations in SI direction are shown in both, the AP and LR view image.

Translation along this axis	Are 'seen' in this view	Appear as pseudo rotations about this axis
AP	LR	SI
SI	LR	AP
	AP	LR
LR	AP	SI

2.10. Testing the Method

To test the algorithm, CT images of the humanoid Pixy phantom were obtained. The isocenter position was chosen inside the torso of the phantom between the hip and arm joints, because controlling the exact positioning of the phantoms moveable arms and legs is cumbersome. Positioning errors of the arms or legs would introduce unintended image dissimilarities if they were visible within the FOV. The CT data set

was imported into the Pinnacle3 ® treatment planning system, where the MDRRs were created as described above in section "2.3. Generation of the MDRR Data Base".

The phantom was carefully positioned on the treatment couch by means of lasers and fiducial markers (1.5 mm ball bearings, Beekley Corp., Bristol, CT) to replicate its position from the CT couch (see also section "2.4. Acquisition of the EPID Image Pairs"). Once the phantom was positioned on the treatment couch, a set of baseline AP and LR view EPID image was collected to establish the initial (baseline) setup position (translations and rotations) of the phantom relative to the CT dataset. To remove any offset caused by slight inconsistency of the phantom position during the CT scan and on the treatment couch, the result of the baseline registration was subtracted from the results of the deliberately translated and rotated cases. The resulting net translations and rotations are free of small setup errors of the phantom on the treatment couch and purely reflect the algorithm's ability to find the purposely introduced setup variations. Typically, one or two monitor units (MU) are used clinically to obtain a patient EPID image. To evaluate the potential clinical utility of this method, all EPID images of the humanoid phantom were acquired using 1 MU (2 MU for the pair) at a dose rate of 300 MU/min.

The method is tested by first measuring its accuracy independently for rotations and translations. These tests are complemented by investigations of the dependence of rotations on translations and vice versa. Lastly, the method's ability to register combined random translations and rotations is investigated.

2.10.1 Translations

Registration of 3D translation was tested by subjecting the Pixy phantom to 20 random translations within an interval of ±20 mm (the setup values). The size of those translations was determined by generating three sets of 20 random numbers, uniformly distributed within an interval of ±20 mm. Each of the 20 triplets of numbers represents a translation in mm along the AP, SI and LR axis, respectively. The phantom was placed on the treatment couch and the combined translations in AP, SI and LR direction realized by moving the couch in vertical, longitudinal and lateral direction respectively. In each position, a pair of orthogonal EPID images was acquired.

The EPID and corresponding MDRR images were then cropped to 10 x 10 cm^2 FOV, pre-processed, registered and the results (found values) for translations compared with the known setup values. Tests were carried out to compare extended image pre-processing with only histogram alignment (basic pre-processing) for each of the three cost functions.

The above approach utilizes couch translations to change the phantom position. This means that the accuracy achieved by our registration method also contains the accuracy of couch translations, which are not known exactly. Furthermore, couch position is displayed to the operator of the linac only in 1 mm increments. This means that the couch position for the un-translated phantom ("zero position") is known only with an accuracy of 1 mm. Since translations were realized by adding the translation distance to the "zero position" value, a systematic shift of up to 1 mm for setup translations is possible.

In an attempt to remove couch positioning errors from the registration results, the following test was carried out: First, a pair of orthogonal EPID images of the un-translated phantom was acquired. Then, the translation of these EPID images was simulated by cropping them to a 10 x 10 cm^2 FOV, which was not centered on the beam axis (i.e. it was translated). The size of those translations was again determined by generating three sets of random numbers, this time 200 for each axis, which were uniformly distributed in an interval of ±15 mm. The AP view EPID image 'sees' phantom translations in SI and LR direction (see also Table 2.1); the second and third number of each triplet are therefore used to translate the AP view EPID image. Since the LR view image 'sees' translations in SI and AP direction, the first and second number in the same triplet are used to translate the LR view image. The images then were again pre-processed and registration accuracy determined for the three axes and cost functions.

2.10.2 Rotations

The ability of the algorithm to find patient rotations was tested by generating three sets of 20 random numbers, each uniformly distributed within an interval of [-4°, +4°]. Each of the 20 triplets of random numbers represents a combination of rotations

(in degrees) of the phantom about the AP, SI and LR axes (rotation vector $\Phi = (\Phi_{AP}, \Phi_{SI}, \Phi_{LR})$). A pair of orthogonal EPID images (AP and LR view) was acquired for each of them by transforming the patient rotation vector into the equivalent couch, gantry and collimator angles of the linac as derived in Appendix A.1. The algorithm was then used to register (find) those known setup angles. It was judged by the differences between the registered (found) and the setup angles. Accuracy (mean differences between registered and setup angles) and precision (standard deviation of those differences) were calculated. The maximum deviation form the mean is also reported to keep track of "worst case" results.

As in any image registration task, the results are influenced by a large number of factors. For this application, these include: CT image acquisition parameters (pixel resolution, slice thickness, pitch and technique factors), MDRR creation parameters (pixel interpolation, pixel resolution and pseudo mega voltage x-ray spectrum) and EPID image acquisition parameters (pixel resolution, frame averaging, bias voltage and acquisition mode, megavoltage beam energy, magnification and MUs). Also important are the image pre-processing steps before registration and the image resolution and FOV used during registration. In the method presented here, the angle interval used for creating the MDRR database (i.e. the density with which the parameter space of rotation angle is sampled) is also crucial. Of those factors, the following were selected to be tested for their influence on the performance of the registration algorithm:

- Image pre-processing: The effect of extended image pre-processing was tested by comparing it to results obtained without the additional steps (basic image pre-processing).
- Parabolic fit: Results with and without applying a parabolic fit around the absolute minimum of the cost function were compared.
- MDRR angular spacing: A larger than the default angular spacing of 0.5° would decrease the number of MDRR images, reducing storage and computing requirements. Therefore the effect of coarser angular spacing on the performance of the registration algorithm was investigated by using an angular spacing of 1.0° and 2.5°. This was

done by terminating the search of the parameter MDRR data base as described in section "2.7.1. Searching the MDRR data base" after the step using 1.0° and 2.5° angular spacing respectively.

- Field-of-view (FOV): A large FOV generally increases the information available for registration, but also raises radiation energy deposited in the patient during EPID image acquisition and results in larger images requiring more memory and computing resources. To test the dependence of the registration results on the FOV, images were cropped to four different square field sizes of 5, 7, 10 and 14 cm^2 while keeping all other registration parameters the same. These field sizes were chosen such that the number of pixels in a larger field EPID image were nearly doubled compared to the immediately preceding field size.

- Pixel size: The native pixel size of EPID images and the pixel size of the MDRR images are generally unequal. The images were therefore re-sampled to identical pixel size for registration purposes. To test the influence of the pixel size on the algorithm, the 10 x 10 cm^2 FOV was re-sampled to 70 x 70, 100 x 100, 140 x 140 and 200 x 200 pixel matrices, again approximately doubling the number of pixels in each step.

- Cost function: The tests mentioned above were each carried out three times, each time using a different cost function.

Values for rotation angles are displayed to the operator of the linac in 0.1° increments. This means that the setup values for rotation angles can potentially display an offset of up to this value, as was the case for translations (where the offset could be up to 1 mm).

2.10.3 Coupling of Rotations and Translations

Recalling the flow chart (Figure 2.1), the algorithm registers translations prior to rotations. After the translations have been registered, the EPID images need to be

back-translated by the found translation values in order to align them with the isocenter of the MDRR images. Because the EPID images have been acquired with a diverging beam, however, they still don't appear like an un-rotated image; rather, there is potential for the algorithm to confuse them as images depicting a rotation (pseudo rotation), as discussed in section "2.9. Coupling of Rotations and Translations". It is important to know if and how significantly the fact that a translation occurred will influence the subsequent registration of rotations.

To test how far the registration results for rotations depend on translations, the phantom was translated along the AP, SI and LR axes separately by 5, 10, 20, and 30 mm, but not rotated. Each time, orthogonal EPID images were acquired at gantry angle 0° and 90° and were then registered with the algorithm. Baseline results of the *un*-translated (and un-rotated) phantom were also acquired and subtracted from the results for the deliberately translated case. The resulting registration values for translations reflect purely the ability of the algorithm to find the setup values (the translations). The registration results for rotations (the found rotation angles, which should ideally be zero) assess to which extent the algorithm registers pseudo rotations due to translations.

Tests were also carried out to determine whether the algorithm assigns pseudo translations in the case that the phantom is not translated but rotated. The answer to this question is obtained by deliberately rotating the phantom separately about the AP, SI and LR axes. The rotation angles used for each axis were 0.5°, 1.0°, 2.0° and 4.0°. Since the phantom was not translated in each of these tests, the expected translation values were zero within the accuracy determined in section "2.10.1 Translations". A systematic dependence of registration results for translations on the setup rotation angles would indicate a coupling of translations to rotations.

2.10.4 Combined Translations and Rotations

In clinical practice rotations and translations do not occur separately; rather a patient's pose is a combination of translations along and rotations about each axis. The method's ability to register combined translations and rotations was tested by introducing 20 random setup translations along each axis. As before (section "2.10.1 Translations"), translations were implemented by moving the treatment couch in

longitudinal, lateral or vertical direction. Each translational setup was then combined with random rotations in the interval of ±4° about each axis, implemented as described in section "2.10.2 Rotations". The orthogonal EPID images acquired for each pose were registered and the registration results compared with the setup values.

In order to evaluate the influence of pseudo rotations on the registration algorithm, this test was repeated twice: The first time for setup translations randomly distributed in the interval of ±15 mm and a second time for the decreased interval span of ±7 mm. Setup values for rotations were identical in both cases.

2.11. References

1 P. Thevenaz, U. E. Ruttimann, and M. Unser, "A pyramid approach to subpixel registration based on intensity", "IEEE Trans. Imag. Proc." **7**, 27-41 (1998).

2 J. T. Bushberg, J. A. Seibert, E. M. Leidholdt, and J. M. Boone, "*The essential physics of medical imaging*", 2nd ed, (Lippincott Williams & Wilkins, 2002).

3 H. E. Johns and J. R. Cunningham, "*The physics of radiology*", 4th ed, (Charles C Thomas, 1983).

4 J. v. Dyk, "*The modern technology of radiation oncology*", (Medical Physics Publishing, 1999).

5 F. M. Khan, "*The physics of radiation therapy*", 3rd ed, (Lippincott Williams & Wilkins, 2003).

6 L. Dong and A. L. Boyer, "An image correlation procedure for digitally reconstructed radiographs and electronic portal images", Int. J. Radiat. Oncol. Biol. Phys. **33**, 1053-1060 (1995).

7 N. Dekker, L. S. Ploeger, and M. van Herk, "Evaluation of cost functions for gray value matching of two-dimensional images in radiotherapy", Med. Phys. **30**, 778-784 (2003).

8 R. C. Gonzales and R. E. Woods, "*Digital image processing*", 2nd ed, (Prentice Hall, 2001).

9 J. V. Hajnal, D. L. G. Hill, and D. J. Hawkes, "*Medical image registration*", (CRC Press, 2001).

10 D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration", Phys. Med. Biol. **46**, R1-R45 (2001).

11 C. E. Shannon, "A mathematical theory of communication", Bell System Technical Journal **27**, 379-423 (1948).

12 W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, "*Numerical recipes in c*", 2 ed, (CAMBRIDGE UNIVERSITY PRESS, Cambridge New York Port Chester Melbourne Sydney, 1992).

13 L. M. Sirois, D. H. Hristov, and B. G. Fallone, "Three-dimensional anatomy setup verification by correlation of orthogonal portal images and digitally reconstructed radiographs", Med. Phys. **26**, 2422-2428 (1999).

Chapter 3

Image Registration: Results and Discussion

3.1. Evaluation Overview

To evaluate the image registration algorithm, a set of known setup values is compared with the corresponding set of values found by the registration method, also called registered or measured values. This comparison is accomplished by calculating the difference between the registered and setup values for each member of the test set. The mean, standard deviation and maximum deviation are calculated from the resulting distribution of differences within the test set. These measures characterize the distribution and allow comparison of results obtained in different tests, for example when changing registration parameters such as image pre-processing steps or cost functions. To gain further insight into the data, the results for the rotation about or translation along each of the three different axes (AP, SI and LR) are shown separately unless stated otherwise. Figure 3.1 illustrates this process of data evaluation for the example of registering rotations and comparing them with the corresponding setup values.

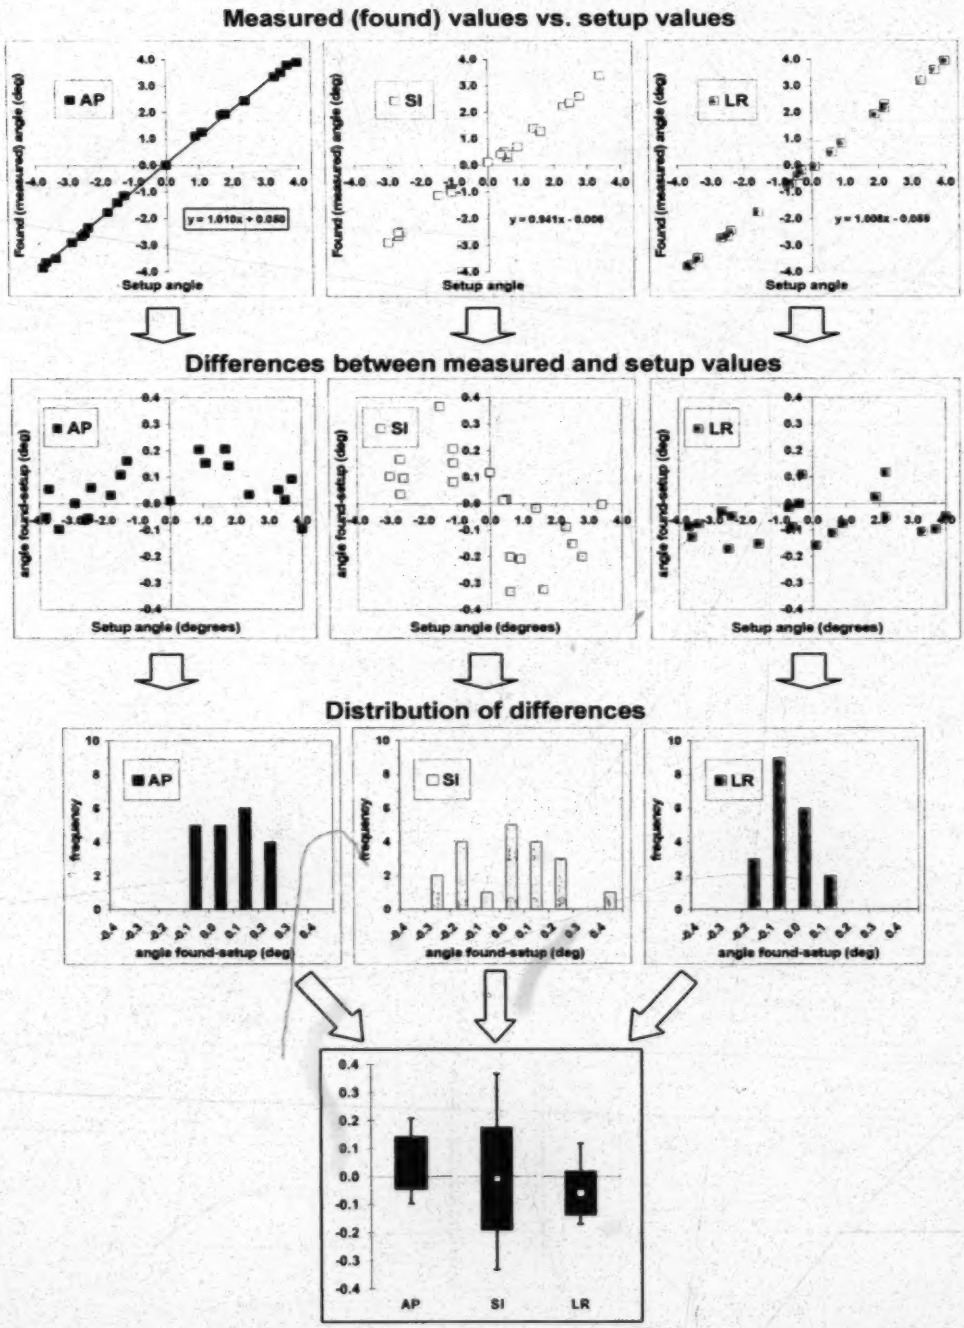

Figure 3.1 Flow chart illustrating how the performance of the registration method is evaluated by comparing registered (found) with setup values, and characterizing the distribution of differences by their average, standard deviation and maximum deviation, corresponding to the dots, bars and whiskers respectively in the bottom chart.

3.2. Translations

Figure 3.2 shows in an array of charts the evaluation of the distribution of differences between the found and the setup values for translations along each axis, grouped by cost function (columns) and the type of image pre-processing (rows).

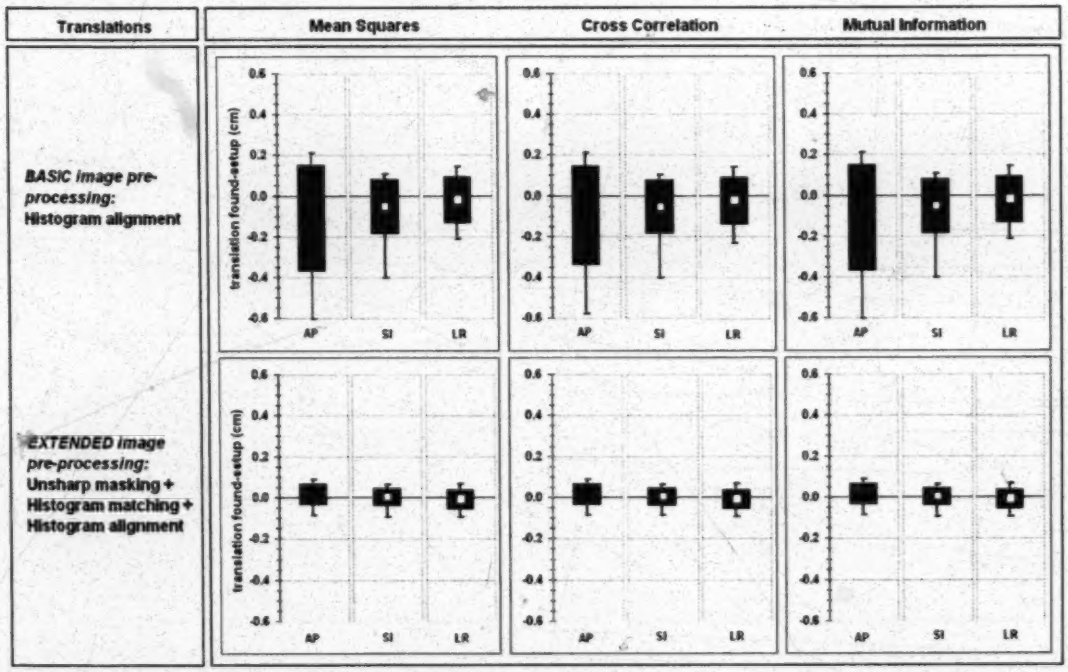

Figure 3.2 Registration results for translations, investigating the effect of using different cost functions and image pre-processing. The charts are evaluating the difference between the translation value found by the algorithm and the random setup values along the AP, LR and SI axes. The average, standard and maximum deviations of the differences are shown as dots, thick error bars and thin error bars respectively.

The use of extended image pre-processing significantly reduced the standard deviation and spread of the distributions. This was observed for all cost functions and axes. Extended image pre-processing also causes the average values to lie closer to zero than basic image pre-processing. Differences between the three cost functions for any given axis were negligible (less than 0.02 mm for average, standard, and maximum deviations for any axis). The first three columns of Table 1 ("Translations of phantom") summarize the results obtained for extended image pre-processing for each axis, while averaging over all cost functions. Registration errors were $\leq$1.0 mm and registration

precision was better than or equal to 0.5 mm (1 SD). It can be noticed, that the standard deviation for the SI translations are smaller than for translations along the AP and LR axes. This is explained by the fact that translations in SI direction are determined twice, namely as the average of values obtained by registering the images in AP-view and the images in LR-view. This additional information available for registration when determining SI translations leads to a slightly more precise result.

Table 3.1 Registering known random translations: Values describe the distribution of differences between setup and found (registered) translations (mm). They are the averages of values obtained for registrations using each of the three cost functions.

	Translations of phantom			Simulated translations		
	AP	SI	LR	AP	SI	LR
Average (mm)	0.2	-0.1	-0.1	0.1	0.0	0.0
Std. dev. (mm)	0.5	0.4	0.5	0.3	0.1	0.1
Max. dev. (mm)	1.0	1.0	1.0	0.4	0.4	0.5

As has been noted in chapter "2.10.1 Translations", the results obtained when translating the phantom with the help of the treatment couch represent the upper limit for the registration accuracy because they also contain the couch positioning errors. Tests carried out with simulated translations excluded the couch positioning error. They resulted in the values listed in the three rightmost columns of Table 1 ("Simulated translations"). Like for actual phantom translations, values obtained for the different cost function were very close and are therefore averaged in Table 1. The values obtained for the standard deviation of simulated translations improve by between 0.2 mm (AP axis) and 0.4 mm (LR axis) compared to actual translations of the phantom; values for maximum deviation are at least halved. The simulated translations of the EPID images however neglect image dissimilarities that are introduced to images of physically translated objects in diverging beam geometry (see Figure 2.16.b). The registration results for simulated translations therefore likely overestimate the accuracy achieved by the registration method. Thus the true value for the registration accuracy lies between the values listed in Table 1: the standard deviation for any axis is greater than 0.3 and less than 0.5 mm; the maximum deviation lies between 0.5 and 1.0 mm.

3.3. Rotations

The influence of image pre-processing and cost function on the ability of the method to register rotations is investigated first. In an array of charts, Figure 3.3 shows the parameters for the distribution of differences between the angles found by the algorithm and the setup angles. Each column of charts displays the results for a different cost function for the three axes. Each row of charts displays the results for different image pre-processing steps applied before registration.

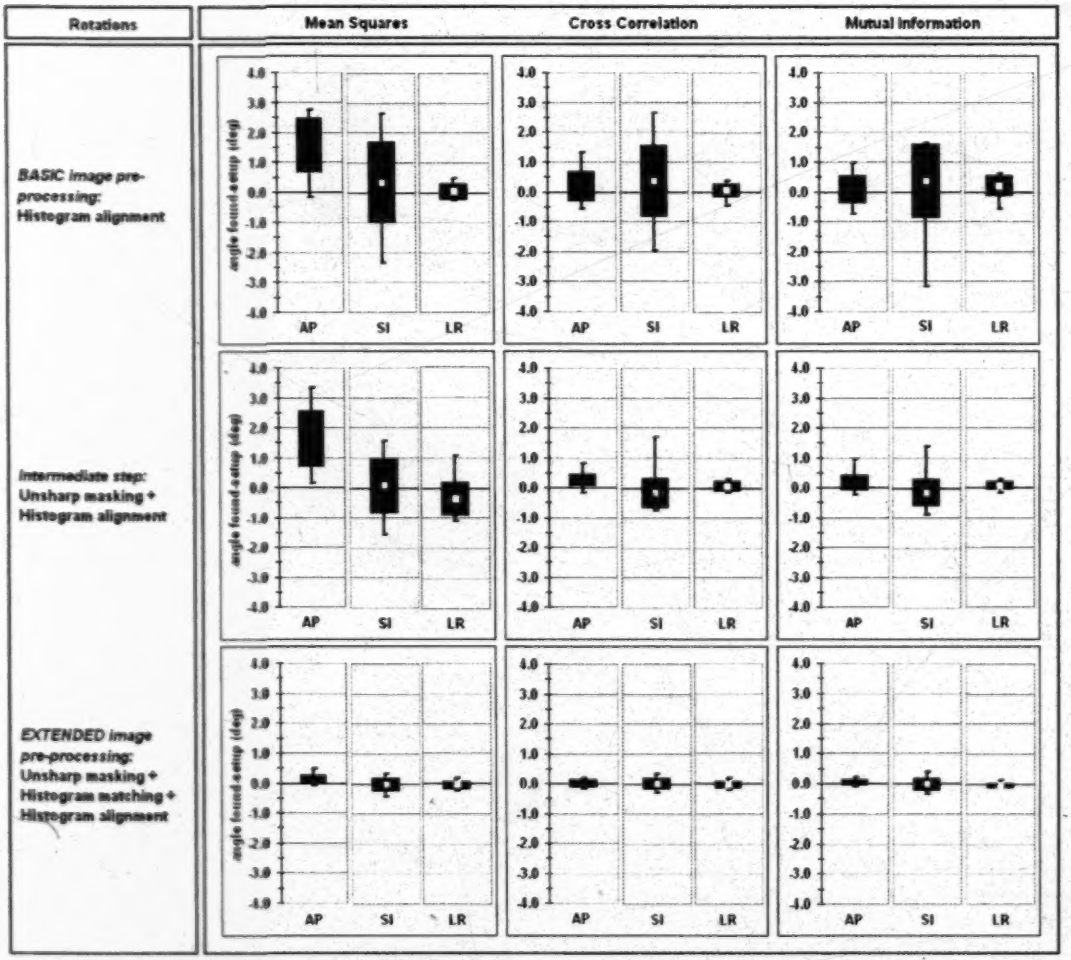

Figure 3.3 Effect of image pre-processing and cost function on registration of rotation angles about each axis. Evaluated are the difference between the angles found by the registration routine and the random, known setup value. Shown is the average value (dots), standard deviation (thick bar) and min and max values (thin bars). Averages close to zero and small bars indicate accurate and precise registration.

The additional image pre-processing steps led to a significant reduction of standard deviation and spread of the results. Compared to the mean squares cost function, cross correlation and mutual information showed smaller standard deviation and closer-to-zero averages, especially for the LR axis or when histogram matching was not performed. Rotations about the SI axis consistently display a broader distribution than rotations about the LR and AP axes. This is likely due to the fact that the rotations about the SI axis do not possess an in-plane component with either the AP or LR view image. The cost function minimum in the parameter space may be less sharp, which would lead to a somewhat less consistent registration result, indicated by a larger standard deviation. The results for the bottom row of charts in Figure 3.3 are summarized in Table 3.2. Mutual information results in the lowest values for standard and maximum deviations for rotations about the AP and LR axes, whereas cross correlation performs best for rotations about the SI axis. In all cases, the standard deviation is less than half the angle interval used to sample the parameter space (0.5°). This is even the case for maximum deviation in the case of cross correlation and mutual information for the AP and LR axes.

Table 3.2 Registering known random rotations: Values represent the differences between angles found by the registration routine and known setup angles (deg). Values are for extended image pre-processing (bottom row of charts in Figure 3.3).

	Mean squared difference			Cross correlation			Mutual Information		
	AP	SI	LR	AP	SI	LR	AP	SI	LR
Average (deg)	0.13	-0.01	-0.05	0.02	0.02	-0.04	0.05	-0.01	-0.06
Std. dev. (deg)	0.15	0.20	0.11	0.12	0.18	0.11	0.09	0.18	0.08
Max. dev. (deg)	0.34	0.42	0.21	0.21	0.34	0.20	0.16	0.38	0.18

The standard deviations in Table 3.2 come close to the positioning accuracy of the linac itself (better than 0.1° for gantry and collimator rotations, better than 0.2° for couch rotation), which means that the observed values could be caused partly by the linac when the EPID images were acquired – especially for rotations around the AP and LR axes in the case of mutual information, where the standard deviation was less than

0.1°. As such, the standard deviations in Table 3.2 represent upper limits for the accuracy of the method.

The values found by the registration algorithm, using extended imaged pre-processing and mutual information, are plotted in Figure 3.4 against the setup values. As was mentioned in chapter 2 "Image Registration: Materials and Methods", the setup angles ranged from -4° and +4°. Ideally, a straight-line fit to the data points would have a slope of unity. The deviations from a slope of 1 for the AP, SI and LR axes were 1.0%, 5.9% and 0.8% respectively. The intercept with the ordinate is less than 0.06° for each axis and smaller than any standard deviation measured (see Table 3.2). All registrations converged, i.e. all angles were found within the interval [-5°, +5°] and all of the eigenvalues of the Hessian matrix (matrix C in Equation 2.6) were positive, indicating that the parabola indeed contains a minimum.

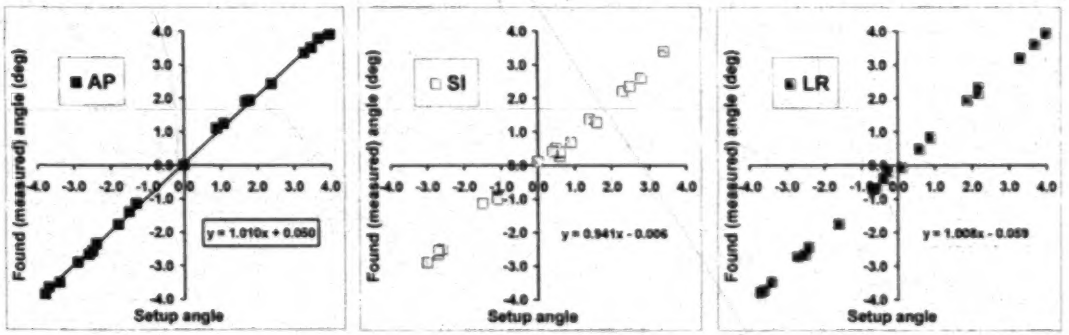

Figure 3.4 Angles found by the registration routine (measured angle) vs. the known setup angles by which the phantom was rotated. Shown are results for each of the three axes (AP, SI, LR) for the cost function based on mutual information. The data has been fitted with a straight line.

The advantage of applying the parabolic fit strategy becomes more apparent when the angular interval between successive MDRRs in the data base is increased (i.e. the parameter space is sampled coarser): angular intervals of 0.5°, 1° and 2.5° were tested. Figure 3.5 compares registration results for three different strategies of finding rotation angles: "Parabolic fit", "MDRR angles" and "From translations" (see Chapter "2.8. Registering 3D Rotations").

All strategies performed similarly for the average value (left chart in Figure 3.5), but when the parabolic fit is applied the values varied much less for different axes and cost function, as indicated by the smaller 'error bar'. The middle chart shows that the standard deviation of misregistration increases for all strategies with increased MDRR angle spacing, but does so much slower when applying the parabolic fit. A similar behavior is observed for the maximum deviation (right chart), where the parabolic fit strategy did not rise until the angular interval was increased beyond 1°.

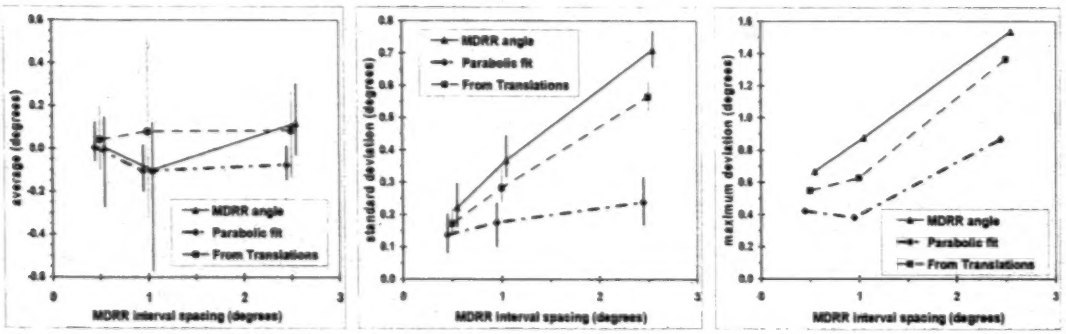

Figure 3.5 Dependence of registration results on angular interval of MDRR data base: Charts evaluate the difference between the setup and found angles of the phantom using three different strategies: "MDRR angle", "Parabolic fit" and "From translations" (each graph's x-axis has been offset slightly to avoid overlap of the data). The left chart shows the average difference, which should be zero for ideal registration. To avoid clutter, values have been averaged for the different cost functions, as well as axes; the 'error bars' depict the range of those values. The middle chart shows the standard deviation of the angle difference. The 'error bars' again indicate the range of values for the different axes and cost functions. The chart on the right shows the maximum deviation of the angle difference from its average value (for all three cost functions and all three axes); it depicts a 'worst case' scenario.

Figure 3.6 shows the frequency distribution of the amount of misregistration observed for each strategy at an angular spacing of 2.5°. While the histogram is peaked around zero for the parabolic fit strategy, it shows approximately a uniform distribution in the interval [-1.25°, +1.25°] for the "MDRR angle" case. The standard deviation for a uniform random distribution is 0.72° (the calculation is performed in Appendix A.2). This value is in very good agreement with the value of 0.71° observed for the "MDRR angle" case". It confirms the intuitive notion that using a MDRR grid with 2.5° angular

spacing, the registration errors for random rotations should be evenly distributed in that 2.5° interval.

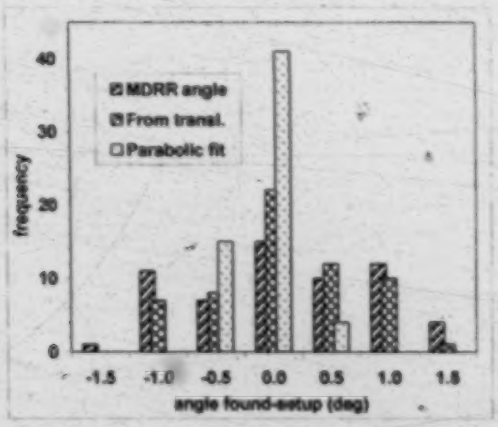

Figure 3.6 Histogram showing the distribution of misregistration (difference between the angle values found and the setup value) for the three different search strategies at an MDRR interval spacing of 2.5°, for all three axes.

In contrast, the "parabolic fit" strategy yielded a standard deviation of only 0.24° when using angular spacing of 2.5° in the MDRR data base, better even than the values obtained for the narrower angular spacing of 1° from the "MDRR angle" or "From translations" strategies (0.28° and 0.37° respectively). Thus, for a given precision, computing and storage requirements can be significantly reduced by employing the parabolic fit. If the angular interval was to be increased from 0.5° to 1°, the data base would shrink by a factor $(21/11)^3 \approx 7$; similarly, increasing it to 2.5 degrees would decrease MDRR computing time and storage requirements by a factor of ≈ 74. Table 3 summarizes these values. It includes the computing times observed and the registration accuracy (largest standard deviation observed, using mutual information).

Clinical implementation of this registration method would be determined by the resources required. Several approaches to reduce the number of necessary MDRRs and to speed up their creation have been published. Sarrut *et al.* [1] suggest decomposing the rotation into an in-plane and out-of-plane component. The in-plane component could be approximated by a 2D rotation of the AP or LR view MDRR, which is not

computer-intensive. Only MDRRs for the out-of-plane rotations would then have to be pre-calculated and stored; or they could possibly, since their number is greatly reduced, be calculated in real-time, which would also allow the use of a gradient strategy to search the angle parameter space instead of using the interval shrinking strategy that results from being restricted to fixed points in the parameter space. Shear-warp factorization [2] is an algorithm approx. 7 times faster than conventional ray-tracing [3] and has the advantage of generating high-quality DRRs. DRR creation can be accelerated even more by first transforming the CT data set into frequency space [4,5]. Ntasis *et al.* [4] extract a specific viewing angle as the inverse Fourier transform of a slice through the center of the frequency space. This method is twice as fast as shear-warp; however, it provides only parallel beam projections. Recently, Rohlfing *et al.* [6] proposed a DRR creation method based on progressive attenuation fields, which is approx. 10 times faster than ray tracing. The influence of those algorithms on the performance of the registration method described here would need to be investigated in future work.

Table 3.3 Registration performance as a function of MDRR angular interval for the "Parabolic fit" strategy (Windows XP workstation equipped with a 2 GHz Xeon processor and 1 GB RAM).

Angular interval	Accuracy of rotations (deg)	MDRR data base size (MB)	Time to create MDRR data base	Time to register one patient position (without image pre-processing, image size: 200x200 pixels)
0.5°	0.06 ± 0.14	2.26 GB	14 hrs	18 min
1.0°	0.14 ± 0.20	333 MB	2 hrs	12 min
2.5°	0.13 ± 0.28	31.3 MB	11 min	5 min

The influence of the field-of-view (FOV) on the registration results was examined by cropping the images to sizes between 5 x 5 cm^2 and 14 x 14 cm^2 while leaving the pixel size unchanged at 0.5 x 0.5 mm^2 as depicted in Figure 3.7. Figure 3.8 shows the results for the different cost functions. The best results were obtained for a 10 x 10 cm^2 FOV. For smaller images, the decrease of image information rendered the registration process less reliable. Larger image sizes also showed an increase in standard and maximum deviation from the average value. Mutual information displayed the smallest standard deviation for each FOV and the slowest rise for larger FOV.

Figure 3.7 Different FOV tested with the routine. Their sizes varied from 5x5 cm^2 to 14x14 cm^2. Shown are the AP (left) and LR (right) view EPID images of the humanoid phantom after pre-processing.

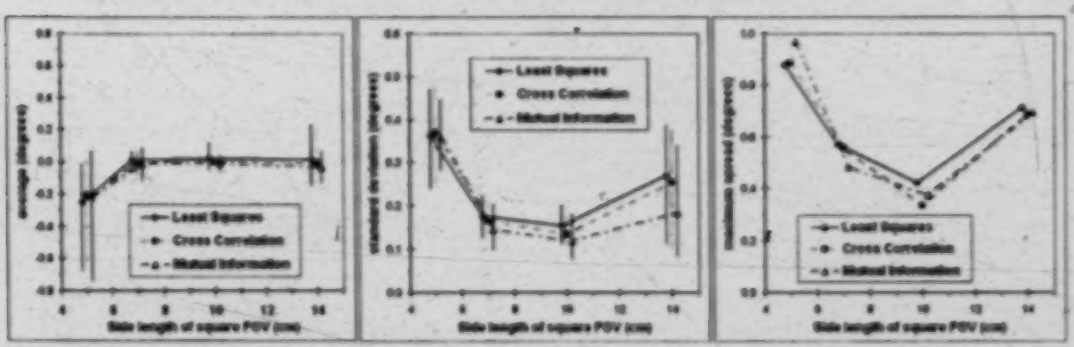

Figure 3.8 Dependence of registration results on the size of the FOV: Charts evaluate the angle difference between the setup and found angles of the phantom. To reduce clutter in the chart, values have been averaged for the three different axes; the 'error bars' depict the range of these values.

While the FOV remained constant at a size of 10 x 10 cm^2, the images were re-sampled to four different pixels sizes; the resulting image thus consisted of 70 x 70, 100 x 100, 140 x 140 or 200 x 200 pixels in order to test the influence of pixel size on the registration method. Figure 3.9 shows the registration results for those configurations. Again, mutual information resulted in the smallest standard and maximum deviation and its average differences were also closest to zero. All three cost functions showed an increase in standard deviation when fewer pixels were used to represent the image. However, using more than 20,000 pixels, corresponding to a re-sampled pixel size of 0.714 mm, did not improve the registration. This pixel size of 0.714 mm in the object plane at 100 cm SAD is magnified to 1.04 mm at the imaging plane since the EPID

imager was located 145 cm from the source. The next smallest pixel size tested was 0.5 mm, which is magnified to 0.725 mm at the imaging plane. The physical pixel size of the EPID imager (0.784 mm) lies between those two values; the EPID's actual spatial resolution is further degraded by its own blur kernel and scatter. Not surprisingly, lowering the sample pixel size below the EPID's spatial resolution did not add to the information content of the image and therefore did not improve registration. However, the re-sampled pixel size of 0.714 mm is less than half the MDRR pixel size of 1.56 mm.

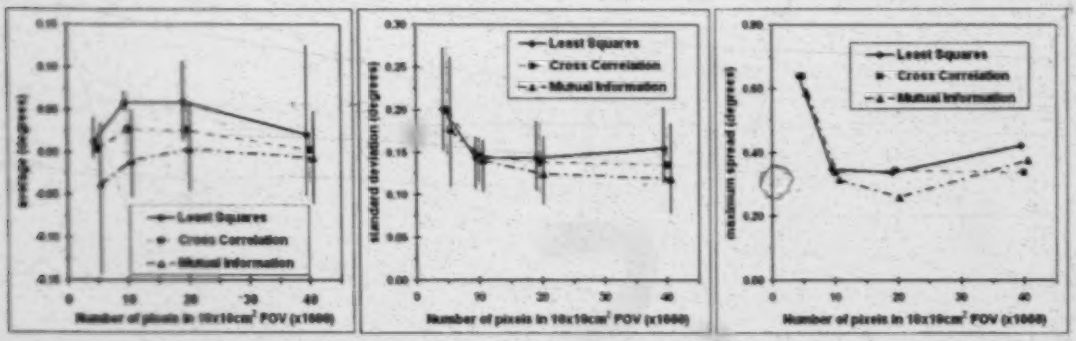

Figure 3.9 Dependence of registration results on the number of pixels in a constant 10 x 10 cm^2 FOV. Charts evaluate the angle difference between the setup and found angles of the phantom. To reduce clutter in the chart, values have been averaged for the three different axes; the 'error bars' depict the range of these values.

In summary, the following combination of parameters provides for best image registration and was used to analyze clinical patient data:

- MDRR angle interval: 0.5°
- Cost function: mutual information
- FOV: 10x10 cm^2
- Re-sampled pixel size: 0.714 mm (140x140 pixels)
- Parabolic fit was applied
- Extended image pre-processing (unsharp masking, histogram matching and histogram alignment)

3.4. Coupling of Rotations and Translations

First, the dependence of registration results for rotations on setup translations was investigated. The phantom was deliberately translated up to 30 mm and the EPID image pairs of the translated phantom were registered. Figure 3.10 shows the registered rotation angles as a function of phantom translation. It can be seen that even though the phantom has not been rotated in these tests, the algorithm finds non-zero (pseudo) rotation angles and their magnitude increases with the amount of translation.

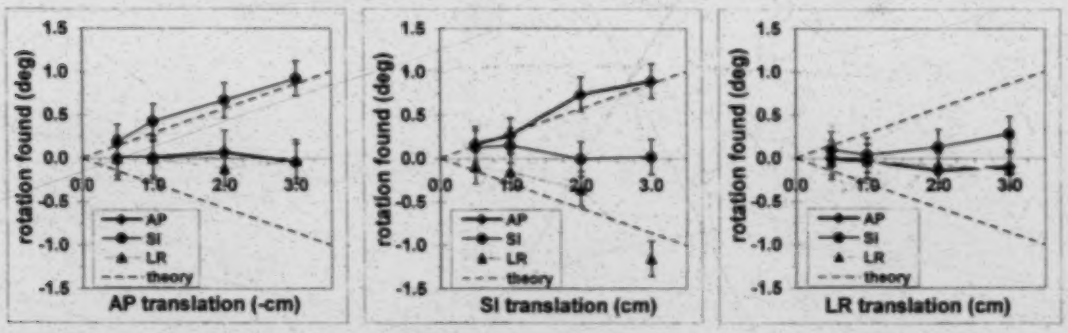

Figure 3.10 Pseudo rotation angles found by the registration method: These occur when the object (in this case the humanoid phantom) is translated because the translated projection image through the object appears similar to an image obtained when the same object has been rotated.

As was discussed in section "2.9 Coupling of Rotations and Translations", a pseudo-rotation is expected to occur as a result of translations *perpendicular* to the beam axis (axis of projection through the data set). This means, that translations along the AP axis are 'seen' by the LR-view EPID images, where they appear as pseudo-rotation about the SI axis that increases with the amount of translation (see Table 2.1). The phenomenon is clearly visible in the left hand chart of Figure 3.10. Translations along the SI axis are 'seen' by both, the AP and LR view EPID images. For the AP view, the SI-translations appear as rotations about the LR axis; for the LR view, those same translations appear as rotations about the AP axis. This is confirmed by the data in the center chart of Figure 3.10, which shows a translation-dependent, apparent phantom rotation about the AP and LR axis. Finally, translations along the LR axis are 'seen' by the AP view EPID image, where they appear as rotations about the SI axis. The onset of this apparent rotation can be seen in the right chart of Figure 3.10, though not as clearly

as in the other two cases. Rotation angles for all other axes remain zero within the error bars, confirming the assumption discussed earlier in chapter "2.9. Coupling of Rotations and Translations" that the magnification changes that occur with translation *parallel* to the beam axis do not influence the registration of rotations.

A theoretical expression for the expected pseudo rotation angles as a result of translations was derived in Equation 2.12 in chapter "2.9. Coupling of Rotations and Translations". The theoretical curves are represented in Figure 3.10 as dashed lines. They agree well within the error bars with the observed values for translations along the AP and SI axis. In the case of translations along the LR axis an apparent rotation about the SI axis is not as large as theoretically expected (right hand chart in Figure 3.10).

The significance of pseudo rotations for the accuracy of registering rotations can be estimated by calculating from Equation (2.12) the distance (translation) at which the pseudo rotations exceed the registration precision for rotations, which was determined to be less than ±0.2° (one standard deviation) for any axis when using mutual information (see Table 3.2). Inverting Equation (2.12) then yields in the case of SOD = 100 cm and $\varphi = 0.2°$: $d = 100\,\text{cm} \cdot \sin(2 \cdot 0.2°) = 0.70\,\text{cm}$. As long as a patient's translational setup errors remain below this value, the registration accuracy for rotations should remain close to the one for the un-translated case.

Next, the dependence of registered translations on setup rotations was investigated by deliberately rotating the phantom and observing the effect of those rotations on the registered translation values. The values found for translations when the phantom was rotated separately about each of the three axes are shown in figure 3.11. Error bars in this chart were taken from the values for "Simulated translations" in Table 3.1 because the phantom had not been translated and therefore linear couch positioning errors are not present in the data.

The registered value for translations agrees with the expected one of zero for most data points within the error bars. Large deviations from zero are observed in the right hand chart for translations along the LR axis. These can be explained as follows: in order to rotate the phantom about the LR axis and acquire EPID images in AP view, the treatment couch needed to be rotated by 90°; this allows the simulation of LR-

rotations by means of setting the gantry angle to the desired LR rotation angle. However, if the treatment couch does not rotate exactly about the isocenter, this couch rotation causes an apparent shift the phantom in SI and/or LR axes when it is viewed in AP direction. This shift appears as translational offset of approx. 1 mm for rotations about the LR axis and has also been documented in monthly quality assurance tests of the couch. It is an artifact caused by the method used to implement known phantom rotations.

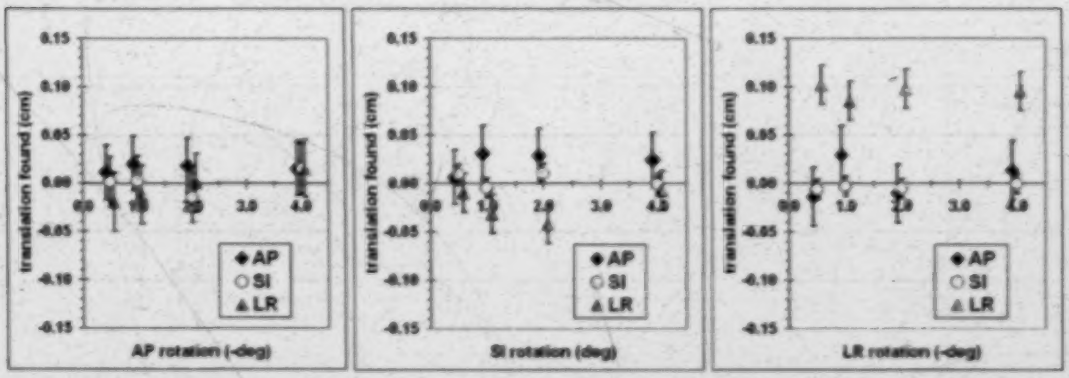

Figure 3.11 Registered values for translations as a function of rotations about the AP, SI and LR axes (left, middle and right chart, respectively) for the un-translated phantom. Error bars for each axis are the averages of values obtained for phantom and electronic translations in Table 1. Each graph's x-axis has been offset slightly to avoid overlap of the data.

In Figure 3.11 the data points for each axis overlap within error bars and therefore show no significant trend for the dependence of registered translations on setup rotations.

3.5. Combined Translations and Rotations

Table 3.4 summarizes the registration results obtained for combined translations and rotations. The largest standard and maximum deviations for translations in the case of the ±7 mm translation interval are 0.5 mm and 1.0 mm respectively. These values agree with the ones obtained for an un-rotated phantom (Table 3.1, columns "Translations of phantom"). An influence of phantom rotations on registration results for translations therefore has not been found if the translations are sufficiently small (≤7

mm in this test). This result also agrees with the findings obtained in the previous section when testing the influence of setup rotations on the registration values for translations for an un-translated phantom (chapter "3.4. Coupling of Rotations and Translations", Figure 3.11); there it was found that the presence of rotations did not significantly affect the registration precision for translations.

The distribution averages for translations are large in the ±7 mm case. As mentioned in Chapter "2.10.1 Translations", this can be explained when considering that setup translations are facilitated by moving the couch relative to the un-translated position. The latter is known only within ±1 mm (the precision with which the couch position is displayed) and thus can create an offset for the distribution of differences.

Table 3.4 Registration performance for combined random translations and rotations. Random rotations were ≤4° about each axis; random translations were within a ±15 mm (top rows) and ±7 mm (bottom rows) interval.

Translation interval		Translations (mm)			Rotations (degrees)		
		AP	SI	LR	AP	SI	LR
±7 mm	Average	-0.6	0.5	0.9	0.11	0.23	-0.13
	Std. dev.	0.4	0.4	0.5	0.19	0.27	0.17
	Max. dev.	1.0	1.0	0.8	0.33	0.50	0.27
±15 mm	Average	-0.1	0.4	0.1	0.12	0.05	-0.20
	Std. dev.	0.6	0.7	0.6	0.37	0.37	0.25
	Max. dev.	1.5	1.6	1.3	0.65	0.91	0.44

Increasing the interval for random setup translations from ±7 mm to ±15 mm, leads to an increase of both, standard and maximum deviations for translations beyond the values found for the un-rotated phantom in Table 3.1. This increase is therefore likely the consequence of the greater image dissimilarity caused by the greater translation, coupled with presence of setup rotations.

When the interval for translations is limited to ±7 mm, standard deviations for rotations are between 0.17° (LR axis), 0.19 ° (AP axis) and 0.27° (SI axis). As previously observed (Chapter "3.3. Rotations"), the SI axis displays the largest standard deviation. Its value is also larger than the standard deviation of 0.18° measured for the

SI axis in the un-translated case (see Table 3.2, "Mutual information"); it documents the onset of the influence of pseudo-rotations on the registration results. This effect increases as the translation interval is widened to ±15 mm. It causes the standard deviation for the registration of rotations to increase further to 0.37° for the SI and AP axes and to 0.25° for the LR axis. Maximum deviations show the same trend. For example the values for the SI axis increase from 0.38° for the un-translated case (see Table 3.2, "Mutual information") to 0.50° if translations up to ±7 mm are introduced and to 0.91° if translations up to ±15 mm are present.

3.6. Summary

Registration accuracy benefits significantly from extended image pre-processing (unsharp masking, histogram matching). The choice of cost function did not influence the registration of translations. Rotations registered most precisely when using mutual information. The registration of rotations also benefited from the interpolation of the parameter space (parabolic fit). Highest precision for registration of three dimensional translations was found to be better than 0.5 mm (one standard deviation) for any axis when no rotations were present. Three-dimensional rotations about any axis were registered with a precision of better than 0.2° (1 SD) when no translations were present. Combined rotations and translations of up to 4.0° and 15 mm were registered with a precision of better than 0.4° and 0.7 mm (1 SD) respectively. This precision improved to 0.3° and 0.5 mm (1 SD) when translations were restricted to ±7 mm.

Setup translations influenced registered values for rotations. The values observed for registered apparent rotations (pseudo rotations) when the un-rotated phantom was translated along the AP and SI axes agree with the theoretical model. Setup rotations did not significantly influence registered values for translations, i.e. pseudo translations were not observed.

3.7. References

1 D. Sarrut and S. Clippe, "Geometrical transformation approximation for 2d/3d intensity-based registration of portal images and ct scan", presented at the Lecture Notes In Computer Science, 4th MICCAI, (2001).

2 P. Lacroute and M. Levoy, "Fast volume rendering using a shear-warp factorization of the viewing transformation", presented at the 21st Conference on Computer Graphics and Interactive Techniques, (1994).

3 J. Weese, R. Goecke, G. P. Penney, P. Desmedt, T. M. Buzug, and H. Schumann, "Fast voxel-based 2d/3d registration algorithm using a volume rendering method based on the shear-warp factorization", presented at the SPIE Conf. Imag. Proc., San Diego, Ca, (1999).

4 E. Ntasis, T. A. Maniatis, and K. S. Nikita, "Fourier volume rendering for real time preview of digital reconstructed radiographs: A web-based implementation", Comp. Med. Imag. Graph. **26**, 1-8 (2002).

5 F. Wang, T. E. Davis, and B. C. Vemuri, "Real-time drr generation using cylindrical harmonics", Lecture Notes in Computer Science **2489**, 671-678 (2002).

6 T. Rohlfing, D. B. Russakoff, J. Denzler, and C. R. Maurer, "Progressive attenuation fields: Fast 2d-3d image registration without precomputation", Medical Image Computing and Computer-Assisted Intervention - Miccai 2004, Pt 1, Proceedings **3216**, 631-638 (2004).

Chapter 4

Biological Impact: Materials and Methods

4.1. Overview

This chapter describes the application of the registration algorithm to clinical treatment plans and the incorporation of calculations of the biological effect of setup errors. The treatment plans were evaluated in retrospect, which meant that only a limited number of orthogonal EPID images were available because treatment protocols currently in use call for orthogonal imaging only once a week. Daily setup verification is based solely on AP view EPID images. The patient setup position was therefore only determined for every 5th treatment fraction in this thesis, six times over the course of the prostate treatments considered here. This smaller number of measurements, however, is only an insignificant limitation and does not limit the applicability of the method. Registration of patient position and the biological impact of setup errors can easily be performed for additional fractions as orthogonal EPID images become available.

Figure 4.1 shows a flow chart, which provides an overview of the process employed in this work to calculate the biological impact of setup errors, once those have been measured (by means of the registration method described in Chapter "2. Image Registration: Materials and Methods"). The original treatment yields a planned

dose distribution, from which DVHs for the regions of interest can be derived. Those DVHs, in turn, serve as the basis for calculating the TCP/NTCP for the original plan. Patient setup errors were included in the dose calculation by first copying the original treatment plan within the Pinnacle 3 ® TPS. Secondly, the patient pose in the copied treatment plan was modified, based on the setup errors (translations and rotations) measured by the registration method. Once the patient pose was modified in the copied treatment plan, the dose distribution for that particular treatment fraction was calculated. This process was repeated for each fraction for which patient position had been measured (i.e. for which orthogonal EPID images were available) and the dose distributions for all fractions were added. The summed dose distribution was displayed in yet another copy of the original treatment plan, referred to as the "sum-plan".

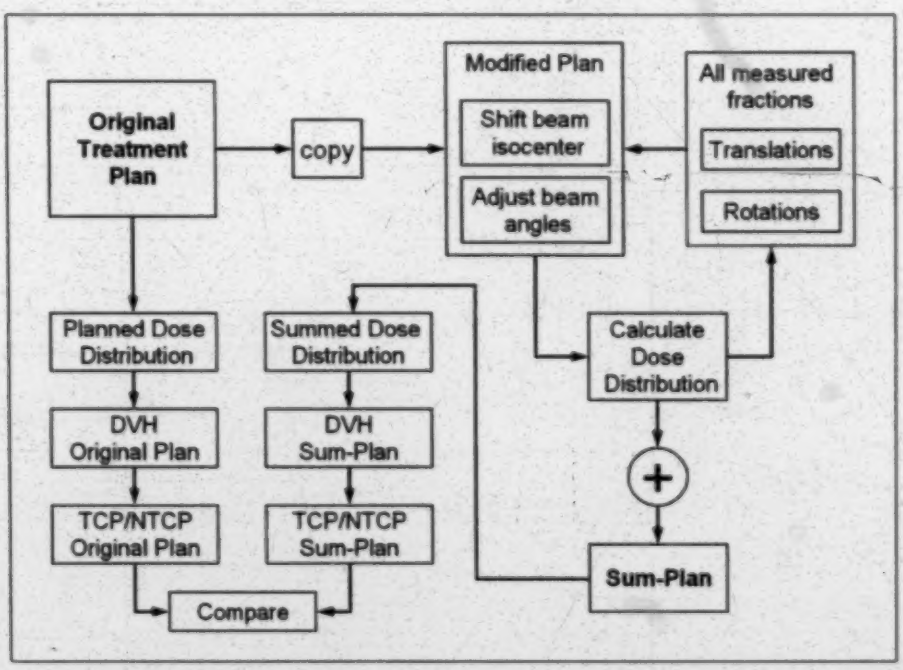

Figure 4.1 Flow chart outlining how patient setup errors (translations and rotations) were processed to evaluate their biological impact on the patient, expressed as TCP/NTCP.

Because the sum-plan was derived as a copy of the original treatment plan, it contained the original contours. Since the dose distribution had changed in the sum plan (reflecting the patient's movement during treatment) relative to the original treatment

plan, so had the dose-volume histograms (DVHs) for the contours. The values for TCP/NTCP are calculated based on the DVHs belonging to a particular contour. Therefore, the TCP/NTCP now also changed and could be compared with the corresponding values obtained for the original plan. The above steps were largely automated and implemented with a C++ plug-in on the Pinnacle3 ® workstation.

4.2. Patient Position

Once the patient positioning error had been determined, it needed to be transferred to and implemented in the Pinnacle TPS, i.e. the original patient position needed to be modified in a copy of the original treatment plan such that it reflects the measured setup error. Because translations and rotations of the patient CT data set within the TPS are not possible, the beam parameters are adjusted such that the actual beams-eye-views (BEVs) of the patient anatomy, i.e. the orientation of the treatment beams relative to the patient, are the same as if the patient data set itself had been moved within the TPS. To that end, patient translations are accounted for by translating the beam isocenter in the opposite direction of the actual (measured) patient translation. Patient rotations are implemented by using the coordinate transform as described in Appendix A.1. The beam angles at which the radiation is delivered in the modified plan are the same as if the CT data set of the patient had been rotated in the TPS.

4.3. Clinical Cases

The algorithm was applied to analyze the treatment received by two patients, both of whom were diagnosed with prostate cancer. The Planning Target Volume (PTV) therefore was a region around the prostate. The critical structure was the rectum, which is sensitive to radiation and is located close to the PTV. The contours of both, the PTV and the rectum were outlined by a radiation oncologist. They served as VOIs for the subsequent calculation of DVHs for each of them. Both patients had finished all fractions of their treatment and orthogonal EPID images were available for 6 fractions for each patient. Table 4.1 summarizes treatment parameters for each of the two patients. Each patient's treatment was divided into two phases: a "main" phase and a

"boost" phase. The latter means that the additional dose was delivered to a reduced PTV by reducing the treatment field size in order to reduce exposure of normal tissue. An example of a main and a boosted field is displayed in Figure 4.2, which shows the BEV of both types of fields.

Table 4.1 Treatment parameters for the two patients

Patient	Plan	Dose delivered (cGy)	Treatment fractions	Orthogonal EPID images available for fractions number
A	Main:	5797	31	1, 6, 17, 27
	Boost:	1496	6	32, 37
B	Main:	5006	31	1, 6, 16, 26
	Boost:	2088	7	32, 38

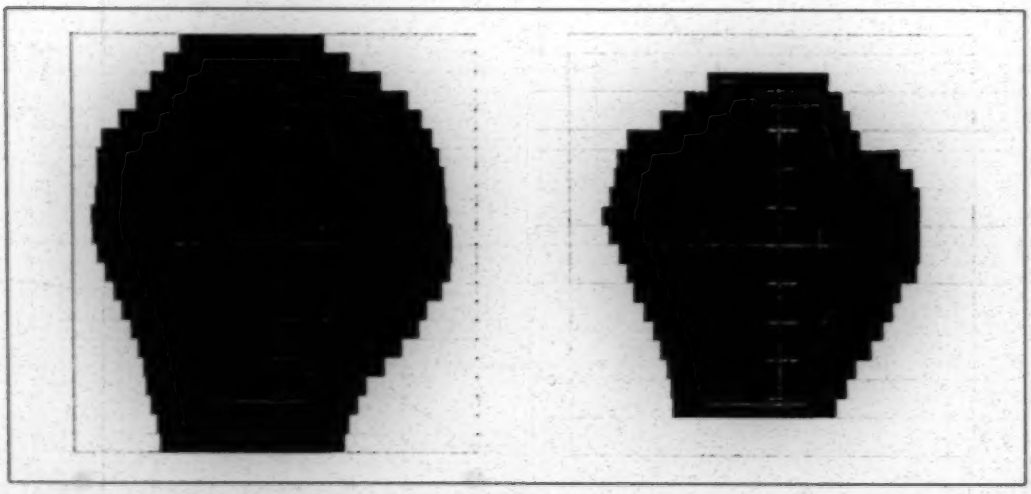

Figure 4.2 BEV of a main field (left) and coned down boost field (right). The beam aperture is shaped with multi-leaf-collimators (MLC, horizontal stripes), which block the beam. The aperture is shaped by the MLC to conform to the shape of the PTV, which contains the prostate (yellow or gray). The width of each MLC leave is 0.5 cm and the height of the PTV is approx. 8.5 cm. The darker (green) structure running along the right side is caused by overlap of the prostate and the rectum (blue).

4.4. Dose Calculation

The treatment for each patient was originally planned on a Helax ® TPS (Nucletron, Veenendaal, The Netherlands). Since the incorporation of setup errors as well as the dose calculation occurs on the Pinnacle 3 ® TPS, the treatment plans needed to be transferred to that work station. This transfer was possible only for the contours of PTV and rectum, but not for the treatment beam parameters (isocenter, treatment angles, field sizes), because of unresolved DICOM (Digital Imaging and Communications in Medicine) connectivity issues between the two TPS. The beam parameters were therefore entered manually into the Pinnacle TPS after the other portions of the treatment plan had been transferred. The position of the MLC leaves was varied manually, starting from a 5 mm margin around the PTV, until the DVHs of the plan generated on Pinnacle for both VOIs (PTV and rectum) closely matched those obtained on the Helax TPS. After generating this "master" treatment plan on Pinnacle (Figure 4.3), it was copied and the beam angles and couch translations in the copied plan were modified to accommodate the measured patient positioning error as described in section "4.2. Patient Position". All dose calculations were performed with the adaptive convolution algorithm, which was part of the Pinnacle TPS.

As was mentioned earlier (section "4.3. Clinical Cases"), the last few fractions of the clinical treatment plans were boosted. Thus, not all of the beams contributed to the whole treatment. The normal beams were only used for the first i fractions, whereas the boost fields were applied for the last j fractions, with $k = i + j$ being the total number of treatment fractions delivered. Whereas the normal and boost fields are treated in separate dose calculations on Helax TPS, they were included in the same treatment plan on the Pinnacle TPS, but as a separate set of treatment beams in order to simplify to dose reconstruction process. This procedure is justified by the following two assumptions:

- The fractions for which patient setup errors were measured are evenly spread out over the whole treatment received by a patient,
- Time trend of setup errors between main (earlier) and boosted (later) fractions are negligible.

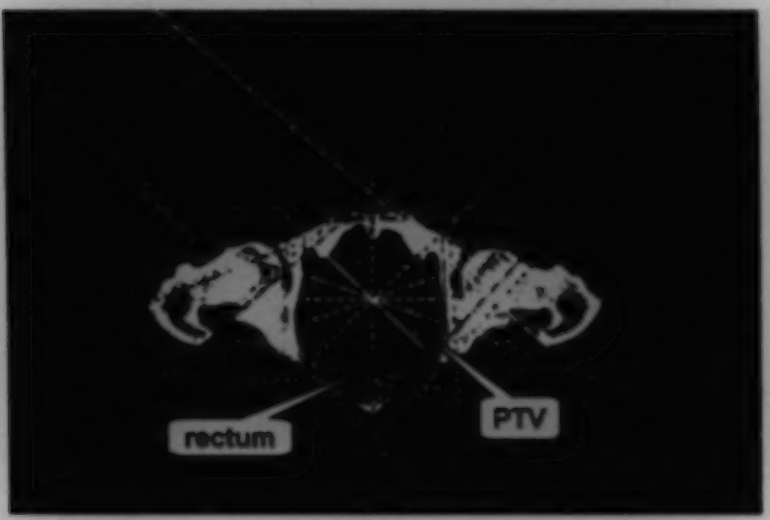

Figure 4.3 Transverse isocenter slice of a "master" treatment plan after it has been set up on Pinnacle. Seven beams were used in this plan.

The first assumption is verified by the fraction numbers for which orthogonal EPID images were available (see rightmost column in Table 4.1). The second assumption is made at this point because no evidence to the contrary exists. It can be verified or falsified once the patient position has actually been measured for each fraction. With those two assumptions it is possible to include both, main and boost fields in each copy of the treatment plan, so that each copy now contained twice the number of beams: half the beams originating from the main and the other half, with a smaller field size but otherwise unchanged beam parameters, from the boost plan. Within a treatment plan, the different beams contributing to the dose delivery are assigned a weight (a number between 0 and 1), which determines how much of the total dose delivered in one fraction is to be delivered by each beam. The weights of all beams add up to one. These weights were assigned in the original (Helax) treatment plan and were also included in the master treatment plan on Pinnacle. However, since main *and* boost field are now present in the same plan, the beam weights of main and boost beams needed to be adjusted in order to add up to one again. The weights of the beams belonging to the main and boost fractions were therefore scaled by a factor i/k and j/k respectively.

Each measured fraction's contribution to the dose in the "sum-plan" was given equal weight. Lastly, the dose values in the resulting dose matrix of the "sum-plan" were divided by the number of measured fractions so that the total number of monitor units delivered to the sum-plan was the same as for the original treatment plan.

4.5. Radiobiological Model

To test the impact of patient setup variations on TCP and NTCP, these values were calculated using radiobiological models which, as mentioned in Chapter "1.4. Biological impact of patient positioning errors", are still being actively developed. For the purpose of this thesis, simple models are used for both, TCP and NTCP. Their choice does not limit the applicability of the dose evaluation process presented in this thesis. They serve not so much to achieve the best accuracy in calculating absolute values of TCP and NTCP, but for the relative comparison of values achieved with different treatment plans. The two models are introduced in the following paragraphs.

To be able to compare both patients, the tumor stage was assumed to be the same for both patients (T1, early stage). As will be seen below, the relationship between TCP/NTCP and dose is a non-linear one. Therefore, in order to allow relative comparison of results for both patients, the initial value for NTCP (for correctly positioned patient) was chosen to be the same. NTCP was set to 5% ± 0.5% for each patient because the value of 5% is the commonly accepted clinical threshold value used in radiation therapy. This scaling was achieved by slightly modifying the total patient dose such that the value for NTCP becomes 5% ± 0.5%. This adjustment was necessary only for patient B where the scale factor was 0.86.

4.5.1. TCP model

It is assumed that a tumor will continue to grow unless all clonogenic cells are killed. The TCP, then, is the probability for zero survival of these cells. Because the number of cells in the tumor is large, leading to low probability of any given cell being hit by a photon, and because the number of photons is large, calculation of TCP lends itself to the application of Poisson statistics if we also assume that damage caused in

one cell does not effect any other cells (independence of events). If the tumor consists of N cells initially, Poisson statistics predicts the likelihood $p(n, D)$ for n surviving cells after administration of a dose D:

$$p(n,D) = \frac{(N \cdot sf(D))^n \cdot e^{-N \cdot sf(D)}}{n!}, \tag{4.1}$$

where $sf(D)$ is the fraction of clonogenic cells surviving the irradiation of the tumor with dose D. To find the TCP, $n = 0$ (no surviving cells) so that:

$$TCP = p(0,D) = e^{-N \cdot sf(D)}. \tag{4.2}$$

The surviving fraction is calculated noting that the rate with which additional volume is damaged by introducing additional dose is inversely proportional to the amount of volume already damaged (neglecting cell repair and repopulation effects):

$$\frac{d\,sf}{d\,D} = -\alpha \cdot sf, \tag{4.3}$$

where α is a constant of proportionality, and therefore:

$$sf(D) = e^{-\alpha D}. \tag{4.4}$$

Inserting Equation 4.4 into Equation 4.2 yields

$$TCP = \exp[-N \cdot \exp(-\alpha D)].$$

This equation has a sigmoidal shape, which can be re-written and parameterized by the values of dose D_{50} and slope γ_{50} for which $TCP = 50\%$:

$$TCP = \left(\frac{1}{2}\right)^{\exp[2\gamma_{50}(1-D/D_{50})/\ln 2]} \tag{4.5}$$

In the case of heterogeneous irradiation, in which different sub-volumes v_i receive different radiation doses D_i and assuming that each sub-volumes' dose response is independent the others', the dose response of the whole tumor is the product of the TCPs for each sub-volume:

$$TCP = \prod_i TCP(D_i, v_i) = \left(\frac{1}{2}\right)^{\sum_i v_i \exp[2\gamma_{50}(1-D_i/D_{50})/\ln 2]} \tag{4.6}$$

This equation is used to fit clinical data of patients' radiation response in order to determine values for D_{50} and γ_{50}, which are then used to predict TCP for new treatment plans. Though the model is initially of a mechanistical nature, its application is phenomenological. The values used in this thesis, which are obtained for prostate cancer with this model are[1]:

$D_{50} = 28.34$ Gy and $\gamma_{50} = 1.0$.

A plot of the TCP curve resulting from these values is shown in Figure 4.4.

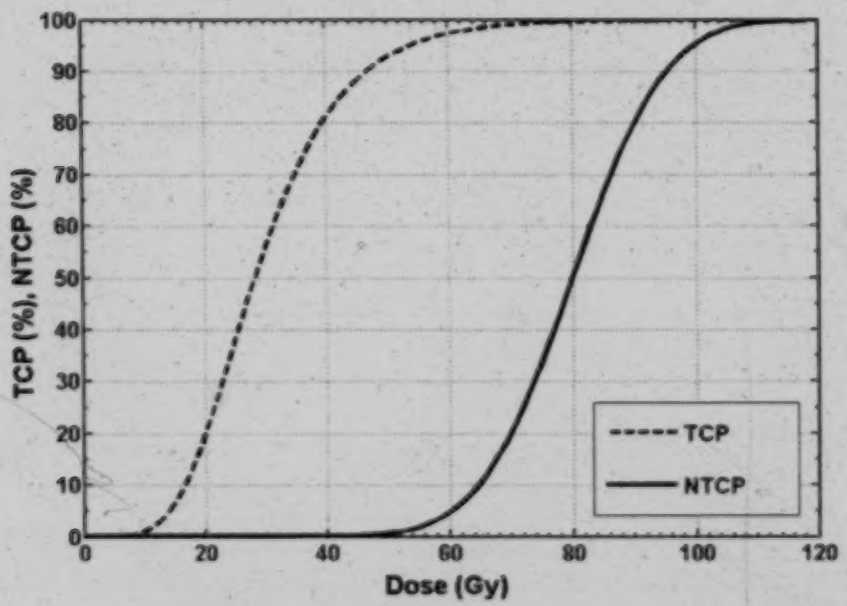

Figure 4.4 Plot of the curves for TCP (Equation 4.6) and NTCP curve (Equation 4.7) for the prostate and rectum, respectively, using the parameters cited in the text.

4.5.2. NTCP model

The Sigmoidal Dose Response (SDR) model, employed in this thesis to describe the response of normal tissue to radiation, has been proposed by Lyman [2] and is largely phenomenological in nature. The sigmoidal shape is realized by means of the error function:

$$NTCP = 0.5 \cdot \left[1 + \mathrm{erf}\left(\frac{x}{\sqrt{2}}\right)\right] = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt, \tag{4.7}$$

with

$$x = \frac{(EUD - D_{50})}{m \cdot D_{50}}. \tag{4.8}$$

Here D_{50} is the dose causing *50%* NTCP and m specifies the slope at this point. *EUD* is the Equivalent Uniform Dose, the dose which, when applied uniformly to the whole organ, would produce the same biological response as the actual (inhomogeneous) dose distribution delivered. In the case of non-uniform irradiation (which is usually the case for any realistic dose distribution), the *EUD* is calculated from the DVH using:

$$EUD = \left(\sum_i v_i D_i^{1/n}\right)^n, \tag{4.9}$$

where the sum is carried out over all sub-volumes v_i of the organ in question, each receiving the dose D_i. *EUD*, m and n represent the three parameters that are used to fit equation 4.7 to clinical dose response data of patients. The organ at risk in the case of this thesis is the rectum, for which the following parameters have been used[3]:

$D_{50} = 80.0$ Gy, $m = 0.15$, $n = 0.12$.

A plot of the NTCP curve resulting from these values is shown in Figure 4.4. Both, TCP and NTCP were calculated with the convenient software tool developed by Warkentin *et al.*[4]

4.6. References

1 P. Okunieff, D. Morgan, A. Niemierko, and H. D. Suit, "Radiation dose-response of human tumors", International Journal of Radiation Oncology Biology Physics **32**, 1227-1237 (1995).

2 J. T. Lyman, "Complication probability as assessed from dose volume histograms", Radiation Research **104**, S13-S19 (1985).

3 C. Burman, G. J. Kutcher, B. Emami, and M. Goitein, "Fitting of normal tissue tolerance data to an analytic-function", International Journal of Radiation Oncology Biology Physics **21**, 123-135 (1991).

4 B. Warkentin, P. Stavrev, N. Stavreva, C. Field, and B. G. Fallone, "A tcp-ntcp estimation module using dvhs and known radiobiological models and parameter sets", J. Appl. Clin. Med. Phys. **5**, 50-63 (2004).

Chapter 5

Biological Impact: Results and Discussion

5.1. Scope

Purpose of this chapter is it to exemplify an application of the method of image registration to clinical data. In a first step, the patients' positioning errors in terms of translations and rotations are determined. In a second step, the resulting changes to the dose distribution as calculated and expressed in the dose volume histograms of two VOIs: the PTV around the prostate and the rectum (OAR). Lastly, the impact of the patient positioning errors on the biological response of the two VOI are determined by computing their TCP and NTCP, respectively, and comparing it to the values obtained for the original treatment plan. It should be noted that these results are meant to provide an example on how the impact of setup errors on biological outcome can be studied. They are not meant to be an exhaustive study of radiobiological clinical outcomes.

5.2. Setup errors

The patients' positioning errors are shown in Figure 5.1 and listed in Table 5.1. The largest systematic translational offset (average translation) exists along the AP axis

of patient A (0.4 cm) and the LR axis of patient B (0.31 cm). The rotational offset is largest around the AP and SI axis of patient A (-1.2° and -1.3° respectively) and around the LR axis for patient B (1.4°). The treatment plan margins of 0.5 cm are exceeded in by translational setup errors in 3 cases for patient A: in the first measured fraction in LR direction and in the 4th and 6th measured fraction in AP direction (see Figure 5.1). Setup variation for patient B remains within the 0.5 cm margin, though just barely for the second measured fraction in AP direction, the 5th measured fraction in LR direction and the 6th measured fraction in SI direction. Since the OAR (rectum) is located posterior to the prostate, special attention is drawn to the translational setup error in AP direction, which is greater for patient A: 0.4 cm compared to 0.15 cm average for patient B (see Table 5.1).

Rotational setup errors are confined to a 'corridor' of roughly ±2.5° (see Figure 5.1). Contrary to translational setup errors, rotations are not explicitly taken into account in the treatment planning process. A rotational setup error of 2.5° would lead to a translation of 0.5 cm for a point at a distance of $0.5cm/\sin(2.5^{\circ}) = 11.6cm$ from the axis of rotation. This distance is slightly larger than the diameter of the treatment beams used in these prostate treatments, which was 11 cm for the normal fractions and 9 cm for the boosted fractions (see Figure 4.2). The translational margins of 0.5 cm that are part of the treatment plan would therefore be sufficient if either the rotational or the translational setup errors occurred without the other. Since both setup errors occur together, the margins might not in all cases be sufficient. The effects of the combination of rotations and translations of course depend on the location in the patient, whether the rotations add to or subtract from the translational setup errors experienced by that location and whether the location is part of either a critical structure or the PTV. The combined effect is best described by actually calculating the dose for each measured fraction and evaluating the DVH for each VOI in question, which is done in the next section.

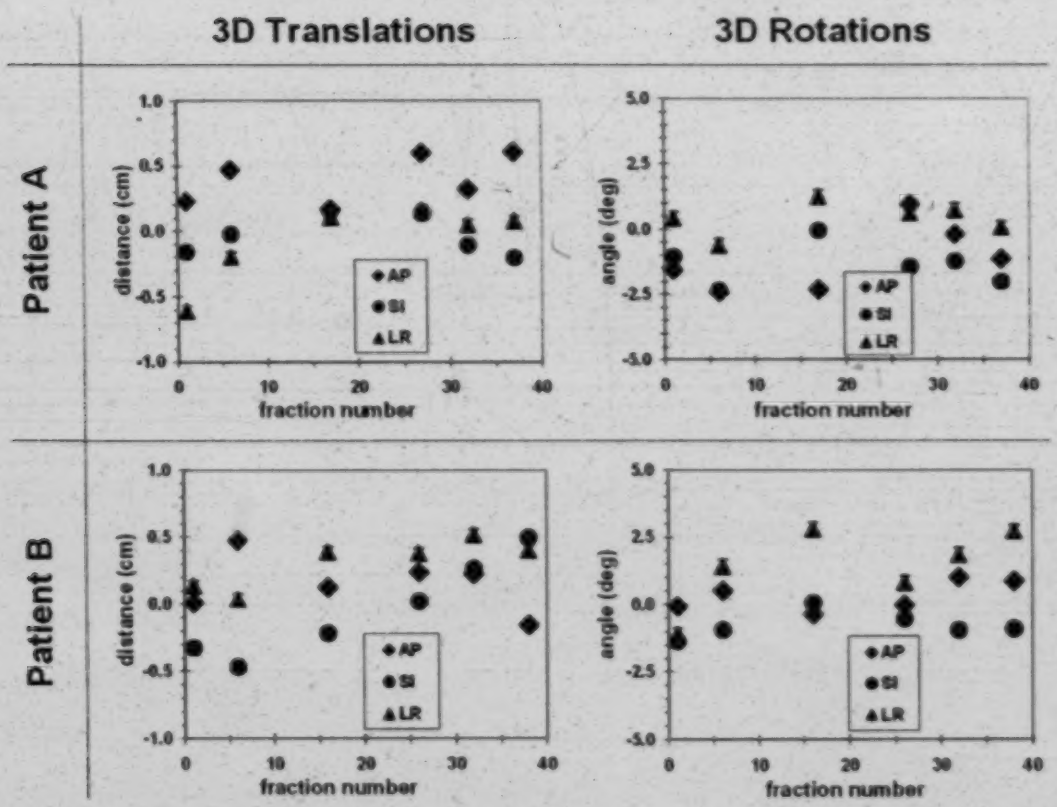

Figure 5.1 Measured setup errors for patient A (top) and B (bottom) over the course of their treatment as indicated by the fraction number on the x-axis

5.3. Dose Distribution and Dose-Volume Histograms

The setup errors were now applied to copies of the original treatment plan and the dose distribution for each fraction calculated and summed up (see flow chart, Figure 4.1). The original and the sum-dose distribution for patient A is shown in Figure 5.2. It has shifted visibly towards the posterior, resulting in an increased dose in the prostate region and large parts of the rectum, which receives an additional dose of approx. 250 cGy to parts of its volume (Figure 5.3). This shift of the dose distribution is consistent with the observed setup errors for Patient A (Figure 5.1), which show a clear translation in AP direction. The summed dose distributions include the measured setup errors and yield the DVH for each of the two VOI (PTV and rectum) of both patients. The DVHs

are displayed in Figure 5.4. In all cases, the DVHs they shifted towards larger dose when setup errors are taken into account. The shift is only minor for the PTV, but pronounced for the rectum for both patients.

Table 5.1 Averages, standard and maximum deviations for clinical patient setup errors, derived from the measurement of fractions per patient shown in Figure 5.1.

		Translations (cm)			Rotations (deg)		
Patient		AP	SI	LR	AP	SI	LR
	Average	0.40	-0.03	-0.07	-1.1	-1.4	0.4
A	Standard dev.	0.19	0.15	0.29	1.3	0.8	0.6
	Maximum dev.	0.23	0.19	0.54	2.1	1.3	1.0
	Average	0.15	-0.04	0.31	0.3	-0.8	1.4
B	Standard dev.	0.22	0.37	0.18	0.6	0.5	1.5
	Maximum dev.	0.32	0.53	0.28	0.7	0.8	2.5

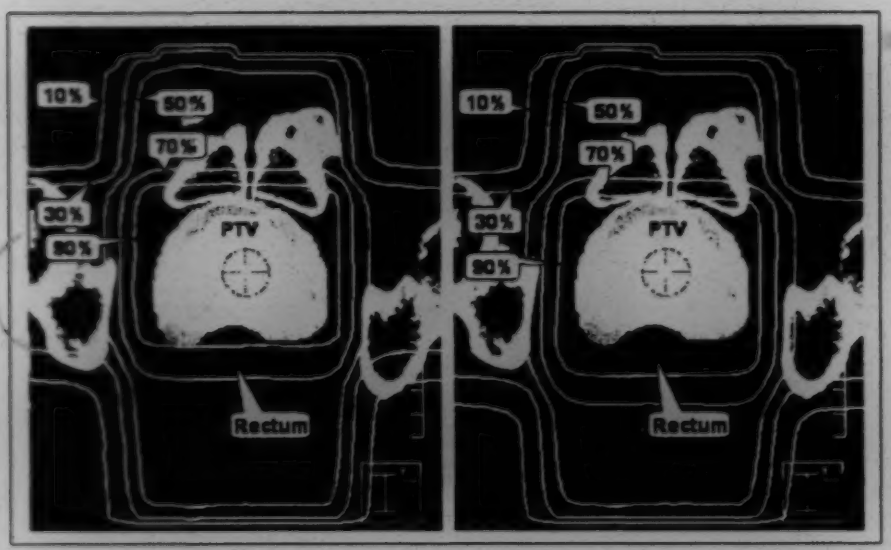

Figure 5.2 Isocentric transverse slice of a prostate treatment plan for patient A, showing the PTV, rectum and isodose lines for the original plan (left) and the summed dose for all setup errors (sum-plan, right). Higher % isodose lines in the sum-plan have moved towards the rectum as a result of setup errors. 100% corresponds to a dose of 72.9 Gy.

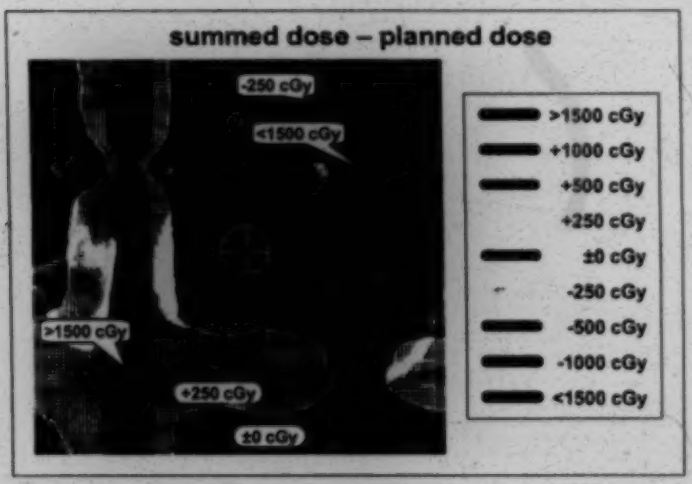

Figure 5.3 Difference of the dose distributions of the sum-plan and the original plan of patient A for the isocentric slice shown in Figure 5.2. An overdose by 250 cGy for large parts of the rectum and a small portion of the prostate is evident.

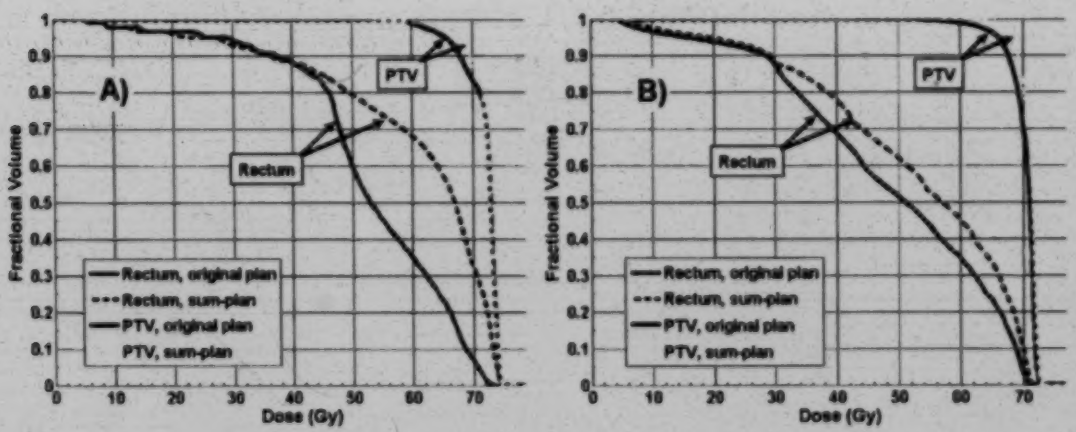

Figure 5.4 Dose-volume-histograms for the PTV and rectum for the original plan and the sum-plan (containing the setup errors) for patient A (left) and B (right).

5.4. TCP and NTCP

The additional dose delivered to the rectum because of patient setup errors, which is evident in Figures 4.2 - 4.3, resulted in an increase to its NTCP. Though present for both patients, this increase was particularly large for patient A (Table 5.2). Since an NTCP of >5% is clinically considered not acceptable[1], the changes in NTCP,

specifically for patient A, would give rise to concern. TCP for the PTV hardly changed for either patient.

Table 5.2 Changes of NTCP and TCP as a result of setup error

Patient	Organ	Original plan	Sum-plan	Change (sum-original)
A	Rectum NTCP	5.24%	14.19%	8.95%
	PTV TCP	99.15%	99.19%	0.04%
B	Rectum NTCP	4.67%	6.92%	2.25%
	PTV TCP	99.01%	99.10%	0.09%

It should be noted that only rigid body transformations of the patient position were assumed, i.e. daily changes of size and position of the prostate and rectum with respect to bony anatomy, which are known to occur,[2-4] are not reflected in these values. This limitation of our method could be overcome only by imaging non-rigid, daily internal anatomy changes using CT or MR scans for each treatment fraction evaluated. This method has been pursued by Schaly *et al.*[5], who retrospectively analyzed CT data sets of five prostate patients. They found that correcting patient setup using portal images resulted in an increase of the minimum dose delivered to the prostate in four of the five patients. Mean dose to the rectum decreased for one and increased for four patients when setup errors based on bony anatomy were corrected.

Happersett *et al.*[6] also evaluated repeat CT images. They found an increase of more than 6% for rectal NTCP in some patients treated with a 6 field boost plan by re-contouring three additional CT scans over the course of the treatment. Those allow access to non-rigid anatomy, but the number of times the re-contouring can be done is limited because of its labor intensity. Considering these differences and the wide spread of NTCP change apparent in even in our very limited sample size of two patients (9.0% and 2.3% respectively), the results are not in contradiction.

5.5. Summary

The potential clinical utility of the registration method developed in this thesis was investigated by applying it to two clinical prostate treatment plans and determining the patient's setup errors. The values for translational and rotational setup variations were then used to re-calculate the patients' dose distributions and dose-volume-histograms as well as changes to TCP and NTCP. For both patients, TCP remained constant, whereas the rectal NTCP increased by 9.0% and 2.2% respectively.

5.6. References

1 J. v. Dyk, "*Modern technology of radiation oncology*", 1 ed, (Medical Physics Publishing, 2005).

2 M. vanHerk, A. Bruce, A. P. G. Kroes, T. Shouman, A. Touw, and J. V. Lebesque, "Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration", International Journal of Radiation Oncology Biology Physics **33**, 1311-1320 (1995).

3 M. S. Hoogeman, M. Van Herk, D. Yan, L. J. Boersma, P. C. M. Koper, and J. V. Lebesque, "A model to simulate day-to-day variations in rectum shape", International Journal of Radiation Oncology Biology Physics **54**, 615-625 (2002).

4 D. Yan, B. Xu, D. Lockman, K. Kota, D. S. Brabbins, J. Wong, and A. A. Martinez, "The influence of interpatient and intrapatient rectum variation on external beam treatment of prostate cancer", International Journal of Radiation Oncology Biology Physics **51**, 1111-1119 (2001).

5 B. Schaly, G. S. Bauman, W. Song, J. J. Battista, and J. Van Dyk, "Dosimetric impact of image-guided 3d conformal radiation therapy of prostate cancer", Phys. Med. Biol. **50**, 3083-3101 (2005).

6 L. Happersett, G. S. Mageras, M. J. Zelefsky, C. M. Burman, S. A. Leibel, C. Chui, Z. Fuks, S. Bull, C. C. Ling, and G. J. Kutcher, "A study of the effects of internal organ motion on dose escalation in conformal prostate treatments", Radiother. Onc. **66**, 263-270 (2003).

Chapter 6

Summary and Conclusion

6.1. Image Registration

A method to register rigid 3D patient position by comparing orthogonal MDRR and EPID images has been implemented and its accuracy has been analyzed as a function of various registration parameters, notably the type of image pre-processing, the cost function used and sampling density of the parameter space of cost function values. Registration precision was better than 0.5 mm (one standard deviation) for 3D translations along any axis if no rotations were present. 3D rotations were registered with a precision of better than 0.2° about any axis if no translations were present. When registering combined rotations and translations the standard deviation for translations remained constant, whereas the one for rotations increased to 0.3° if the magnitude of setup translations was $\leq$7 mm. For setup translations of up to $\pm$15 mm, the standard deviations increased for both, translations and rotations to the values of to 0.7 mm and 0.4° respectively.

The high accuracies achieved meant that the small degree to which setup translations influence the registration results and vice versa become noticeable. The assumption that translations can be registered independently from rotations, which is

strictly valid only for parallel beam geometry, was tested by investigating the influence of setup translations on rotations and *vice versa*. A simple theoretical model for the dependence of rotations on setup translations has been introduced. This model agrees with (AP and SI axes) or overestimates (LR axis) the dependence of registration results for rotation on setup translations (Figure 3.10). As a consequence of this coupling, the registration precision for rotations decreases with increasing translation of the phantom or patient (Table 3.4). A significant dependence of registration results for translations on setup rotations has not been found for the un-translated phantom (Figure 3.11). If, however, rotations and translations were combined, standard deviations for translations did increase if the setup translations were within a ±15 mm interval. No such increase was observed if the setup translations were within a ±7 mm interval (Table 3.4).

It should be noted, that the coupling of out-of-plane rotations to setup translations is a phenomenon that can influence registration results wherever 2D-3D registration is performed with a diverging beam and sufficient accuracy. Once the in-plane translations are known, their effect on out-of-plane rotations can be calculated (Equation 2.11) and corrected for if only one 2D image is used (for example in fluoroscopic imaging[1]). The coupling can be weaker by a factor of two if, as is the case in this thesis, two 2D images are used for registration (Equation 2.12). However, correcting the registration result for pseudo rotations is not as straight-forward in this case, because translations in AP direction ('seen' by the LR-view EPID image) and LR direction ('seen' by the AP-view EPID image) both cause pseudo rotations about the SI axis. Their contributions can either add up or neutralize each other, depending on the direction of each translation.

Image pre-processing significantly improved registration of translations and had an even larger impact on registering rotations. Rotations also benefited strongly from interpolating the space of cost-function values with a parabolic fit. Angular spacing of the MDRR images is important, but increasing from 0.5° to 1.0° only marginally weakens the registration; however it does offer significant benefits for storage and computing requirements and is therefore likely to be preferred in a clinical setting. Clinical implementation could be further aided by reducing storage requirement

and MDRR computing times through more efficient algorithms, the effects of which onto registration accuracy would need to be investigated.

Though in some cases the different cost functions performed similarly, mutual information showed the best overall results, followed by cross-correlation and mean squared difference. This is in contradiction with Penney *et al.*[2], who tested six cost functions and found that mutual information performs worst when registering fluoroscopy images to CT data sets. However, Penney *et al.* did not use *normalized* mutual information, which has been shown to perform more robustly[3] and image pre-processing was not employed. The algorithm proposed by Sarrut *et al.*[4] (decomposing rotations into in- and out-of-plane components) had been tested in reference [5] for six positions of a humanoid Alderson Rando phantom; out-of-plane rotations were sampled every 0.5°. Differences between setup and found translation range up to 2.2 mm and differences between setup and found angle reach up to 2.8°, which is approx. 5 times larger than our values. But their registration code was faster (approx. 3 minutes for one registration) and their data base requirements were reduced (900 DRR images).

The method investigated here registers rigid data sets. As such, a clinical implementation would benefit most where knowledge of patient position, excluding internal organ motion, is desired or for sites where internal anatomy is not expected to change significantly between fractions, such as brain and head/neck. However, the aspects of image pre-processing discussed here might also be useful for situations where non-rigid registration algorithms are employed and/or 3D data sets are registered.

The individual registrations undertaken in this thesis did all converge. This is due to the fact that patient rotations and translation were confined to intervals that are likely to contain clinically relevant setup variations: [-4°, +4°] for rotations and [-2 cm, 2 cm] for translations. If larger intervals are used, the method would perform less stable. This has been observed in preliminary investigations of larger translations (>3 cm), but has not been quantified in detail because it is not relevant for the clinical applicability of the method presented and therefore falls out of the scope of this thesis.

6.2. Biological Impact

A method to study the biological impact of patient setup variations has been presented. The method is based on modifying the dose distribution delivered to the patient such that it contains the measured patient positioning error. The cumulative effect of multiple treatment fractions (if patient position has been measured by acquisition of orthogonal EPID images) can be evaluated with this method. To this end, dose distributions for the measured treatment fractions were added. Two patient treatment plans were investigated and it was found for both that the summed dose distribution changed significantly when compared to the original plan. This change resulted in the increase of NTCP values for both patients above clinically acceptable levels by 9.0% and 2.3% respectively. TCP values on the other hand remained almost unchanged. These results agree with the findings of Happersett *et al.* [6], who also found an increase of rectal NTCP (6%) as a result of setup errors. It should be noted that the sample size (2) in these tests is limited. The data presented therefore illustrates in two examples the applicability of our method to clinical cases. However, the sample size is too small to use the results to characterize in general the clinical radiobiological effect of setup errors.

6.3. Future Developments

Since the dose needed to acquire EPID images is small, it is possible to acquire daily images and monitor a patient's dose distribution including setup errors as the treatment progresses through all fractions. The treatment plan could then be adapted in later fractions, based on measurements done in earlier fractions. This adaptation could be based either on the measured systematic setup error to improve patient alignment[7] or on the re-computed dose distribution to make up for cold spots resulting from setup errors[8]. Ideally, though, the method should be accelerated enough to allow for real-time registration in each treatment fraction between image acquisition and dose delivery. The time for one registration varied between 5 and 18 minutes on a 2 GHz Xeon processor, depending on the angular interval chosen (i.e. registration accuracy, see Table 3.3) because the MDRR data base did not need to be re-calculated for each registration. If

the registration method is to be used on-line to correct patient setup errors, registration needs to be completed within a clinically acceptable time limit of approx. 1 minute. Accelerating the registration by a factor of at least 5, ideally 20, is therefore desirable. The algorithm could be further accelerated by a number of means: changing the functions currently implemented in Matlab ® to a lower level computing language, using a faster work station, employing a more efficient search strategy of the parameter space (for example a form of gradient search) coupled with a fast MDRR calculation algorithm and validating the separate registration of in-plane and out-of-plane components of 3D rotations (thus mostly eliminating the need for MDRRs created for in-plane rotations). The couch translations and the couch, gantry and collimator rotations could then be adjusted on-line to reduce the random portion of the setup errors, thus optimizing dose delivery by adapting the treatment parameters in each fraction.

It has been demonstrated that image pre-processing is significant for image registration. Therefore the influence of image pre-processing on the performance of other image registration problems should be investigated more closely. Obvious candidates would be the inter- and intra-modality registration of 3D data sets generated in CT, MRI, PET and Tomotherapy studies.

A simple geometrical model, describing how translations can appear as pseudo rotations in the registration, was developed. It should be investigated how far this knowledge about pseudo-rotations can be used to correct the rotation values obtained with our algorithm. Care must be taken because pseudo rotations resulting from translations in the positive LR axis (seen by the AP view) and AP axis (seen by the LR view) can cancel each other in the current form of our algorithm. However, some measure of correction should be possible, at least for translations along the SI axis. It should be noted that these pseudo rotations would also appear in registration tasks in fluoroscopic imaging.[1] If out-of-plane rotations are registered, a correction for pseudo rotations could be applied.

The feasibility of the method for inter-fractional biological impact of patient setup variations has been demonstrated. Since it has also been shown that biological

impact can be significant, the method employed in this thesis can provide the groundwork for future investigations with larger patient populations. Those investigations could for example study the effect of different safety margins used when designing treatment plans or the influence of different immobilization devices on the treatment outcome. Since it provides a means to evaluate patient treatment more realistically, clinical studies of treatment outcome could incorporate this information to allow for more accurate development of biological models describing the response of tissue to radiation.

6.4. References

1 J. Weese, T. M. Buzug, C. Lorenz, and C. Fassnacht, "An approach to 2d/3d registration of a vertebra in 2d x-ray fluoroscopies with 3d ct images", in *Cvrmed-mrcas'97* (1997), Vol. 1205, pp. 119-128.

2 G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2-d-3-d medical image registration", IEEE Trans. Med. Imag. **17**, 586-595 (1998).

3 C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3d medical image alignment", Patt. Recog. **32**, 71-86 (1999).

4 D. Sarrut and S. Clippe, "Geometrical transformation approximation for 2d/3d intensity-based registration of portal images and ct scan", presented at the Lecture Notes In Computer Science, 4th MICCAI, (2001).

5 S. Clippe, D. Sarrut, C. Malet, S. Miguet, C. Ginestet, and C. Carrie, "Patient setup error measurement using 3d intensity-based image registration techniques", Int. J. Radiat. Oncol. Biol. Phys. **56**, 259-265 (2003).

6 L. Happersett, G. S. Mageras, M. J. Zelefsky, C. M. Burman, S. A. Leibel, C. Chui, Z. Fuks, S. Bull, C. C. Ling, and G. J. Kutcher, "A study of the effects of internal organ motion on dose escalation in conformal prostate treatments", Radiother. Onc. **66**, 263-270 (2003).

7 T. Bortfeld, M. van Herk, and S. B. Jiang, "When should systematic patient positioning errors in radiotherapy be corrected?" Phys. Med. Biol. **47**, N297-302 (2002).

8 H. Rehbinder, C. Forsgren, and J. Lof, "Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery", Med. Phys. **31**, 3363-3371 (2004).

Appendix

A.1. Patient and Linac rotation angles

A patient's rotation about his AP, SI and LR axis can be equated with corresponding orientations of the linac's beam source, defined by its couch, gantry and collimator angles. The need to determine this correspondence arises several times within this thesis. In a first scenario, patient rotations are simulated by adjusting linac angles (couch, gantry and collimator angles) instead of rotating the patient or phantom. Here, the values to which a given set of linac angles has to be changed in order to afford the same beam's eye view (BEV) of the patient (or phantom) as if the patient was rotated a given amount about his AP, SI and LR axes is calculated. This situation occurs in three sections of this thesis:

1. When the MDRR data base is to be created, the TPS does not allow for the direct rotation of the patient (the CT data set). Instead, the linac's couch gantry and collimator angles are adjusted as to afford an equivalent BEV onto the patient.
2. When EPID images of the rotated Pixy phantom are to be acquired, precise rotations of the life-size phantom are impractical. Therefore, the phantom rotation is again facilitated by adjusting the linac's couch, gantry

and collimator angles to values that result in an equivalent BEV of the patient.

3. When adjusting treatment fields in order to re-calculate the dose received by a patient whose rotations have been found (registered) with our algorithm. Because the TPS does not allow direct rotation of the patient's CT data set, treatment beam angles have to be calculated that again afford the equivalent beams eye view.

A second, related scenario arises when actual patient rotations are known and the linac's couch, gantry and collimator angles are to be changed in order to cancel the misalignment caused by the patient's rotational setup errors. In this case, the amount of patient rotation about his AP, SI and LR axes has been measured and the values to which the planned linac angles (couch, gantry, collimator) need to be adjusted in order to irradiate the patient in the same way as if he had been positioned correctly (i.e. un-rotated) are calculated. Once our algorithm is accelerated enough, this registration could be performed "on-line" in the time between the acquisition of the orthogonal EPID images and the delivery of the treatment. Setup errors can then be corrected by adjusting linac angles to the new values (as a form of Adaptive Radiation Therapy (ART)) before the treatment commences, thus improving the accuracy of dose delivery.

A.1.1 Coordinate transform

Both of the above mentioned scenarios are treated by referencing a fixed (room) coordinate system, relative to which the patient and the linac rotate (Figure A.1). The linac's coordinate system is called "BEV" system. The treatment isocenter is selected as the origin around which all rotations occur.

The linac's couch, gantry and collimator angles relative to the room coordinate system enter into a transformation M. These angles are set by the operator during treatment. The linac angles as they appear relative to the patient determine the transformation $\hat{M}$. The patient's rotation angles around his AP, SI and LR axis form the solid body rotation R. Each of the transformations M, $\hat{M}$ and R can be expressed

as rotational matrices and the linac's couch, gantry and collimator angles can be extracted from M (see below).

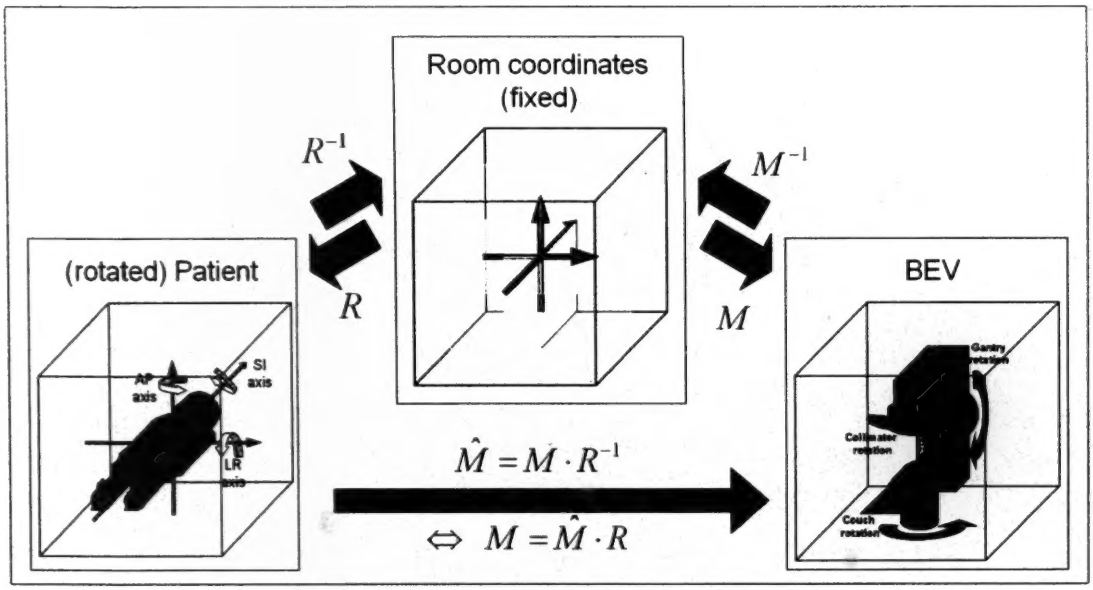

Figure A.1 Coordinate transforms between the Patient, BEV and Room coordinate systems.

Both scenarios described in the previous section ask for the transformation matrix M (containing the linac angles to be set by the operator). As illustrated in Figure A.1, we can determine the transformation M as follows:

$$\hat{M} = M \cdot R^{-1} \quad \Leftrightarrow \quad M = \hat{M} \cdot R. \tag{A1}$$

Obtaining the linac angles for the second scenario, in which linac angles are adjusted in order to compensate patient rotation is now straightforward: the desired (planned) linac angles for each beam are used to calculate $\hat{M}$ in its matrix representation. $\hat{M}$ is then multiplied from right with the matrix R representing the patient's rotation and the resulting matrix M contains the linac angles to be set by the operator.

In the case of the first scenario, which has more relevance for this thesis, we want to simulate patient or phantom rotations R by adjusting the linac's angles instead of rotating the patient or phantom itself. This is accomplished by entering the inverse R^{-1} of R instead of R into equation A1:

$$M = \hat{M} \cdot R^{-1}. \tag{A2}$$

The linac will then 'compensate' the rotation R^{-1}, which, since the patient (phantom) has not actually been rotated, affords the same BEV of the patient (phantom) as if they had been rotated by R and the linac angles had remained unchanged. A simple example will clarify this scenario. Let the gantry angle of the beam to be delivered be 45° relative to the patient, who is located on the treatment couch in a head-first supine (HFS) position. If a patient rotation of +10° about the SI axis was to be simulated, the linac's gantry angle would need to be adjusted by -10° (in the opposite direction of the patient rotation) and thus set to 35° while the patient is kept und-rotated, in order to achieve the same beams-eye-view as if the patient itself had been rotated.

A.1.2 Extracting linac angles from the transformation matrix M

Let R_{AP}, R_{SI} and R_{LR} be the rotation matrices for rotations around the AP, SI and LR axes respectively; then

$$R = R_{LR} \cdot R_{SI} \cdot R_{AP} \tag{A3}$$

describes the 3D patient rotation matrix relative to the fixed (room) coordinate system.

The beams-eye-view of the patient is defined by the collimator coordinate system of the linear accelerator. As discussed in Ref.[1] the room coordinate system is transformed into the collimator coordinate system by

$$M = M_{Col} \cdot M_{Gan} \cdot M_{Cou}^{-1}, \tag{A4}$$

where M_{Col}, M_{Gan} and M_{Cou} are rotational matrices for the collimator, gantry and couch respectively:

$$M_{Cou} = M_{Cou}(\vartheta) = \begin{pmatrix} \cos\vartheta & \sin\vartheta & 0 \\ -\sin\vartheta & \cos\vartheta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{A5a}$$

$$M_{Gan} = M_{Gan}(\varphi) = \begin{pmatrix} \cos\varphi & 0 & \sin\varphi \\ 0 & 1 & 0 \\ -\sin\varphi & 0 & \cos\varphi \end{pmatrix}, \tag{A5b}$$

$$M_{Col} = M_{Col}(\psi) = \begin{pmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{A5c}$$

and

$$M = M(\psi, \vartheta, \varphi) = \begin{pmatrix} \sin\psi\sin\vartheta + \cos\psi\cos\varphi\cos\vartheta & \sin\psi\cos\vartheta - \cos\psi\cos\varphi\sin\vartheta & \cos\psi\sin\varphi \\ \cos\psi\sin\vartheta - \sin\psi\cos\varphi\cos\vartheta & \cos\psi\cos\vartheta + \sin\psi\cos\varphi\sin\vartheta & -\sin\psi\sin\varphi \\ -\sin\varphi\cos\vartheta & \sin\varphi\sin\vartheta & \cos\varphi \end{pmatrix}, \tag{A6}$$

where ϑ , φ and ψ are the couch, gantry and collimator angle, respectively (since we are only interested in the rotational component of the coordinate transform, we have, unlike Starkschall[1], neglected the translational component of the transform into the linac head by the source-to-axis distance (SAD)).

Let $\hat{M}$ describe a specific, desired linac configuration, for example $\vartheta = 0°$, $\varphi = 90°$ and $\psi = 0°$ for a LR view image; then $\hat{M} = M_{Col}(0°) \cdot M_{Gan}(90°) \cdot M_{Cou}(0°)$. If the patient is not rotated (i.e. $R = 1$ (identity matrix)), then set rotations of the linac are trivially $M = \hat{M} \cdot 1 = \hat{M}$. If the simulated rotation of the patient is given by $R \neq 1$, and the patient is placed un-rotated on the treatment couch, then the set rotations of the linac in for the desired beams-eye-view are obtained by Equation A2.

Once the matrix M (equation A6) has been calculated using either equation A1 or A2 (depending on the problem to be solved) the couch gantry and collimator angles are extracted from it:[2]

$$\varphi = \cos^{-1}(m_{3,3}), \tag{A7a}$$

$$\vartheta = \tan^{-1}\left(-m_{2,3} / m_{1,3}\right), \tag{A7b}$$

$$\psi = \tan^{-1}\left(-m_{3,2} / m_{3,1}\right), \tag{A7c}$$

where $m_{i,k}$ are the matrix elements of $M = \widetilde{M} \cdot R^{-1}$ in the i^{th} column and k^{th} row as defined in Equation A6. If $m_{1,3} = 0$ and $\varphi \neq 0°$ in Equation A7b, then $\vartheta = 90°$. If $m_{3,1} = 0$ and $\varphi \neq 0°$ in Equation A7c, then $\psi = 90°$. If $\varphi = 0°$ (gantry vertical), $m_{1,1}$ in Equation A6 yields:

$$m_{1,1} = \sin\psi \sin\vartheta + \cos\psi \cos\vartheta = \cos(\psi - \vartheta), \qquad \text{(A8)}$$

i. e. rotating the couch in one direction or the collimator in the opposite direction yields the same patient orientation, which is obviously the case for vertical gantry. In that case, the couch angle was arbitrarily set to zero and the patient orientation facilitated with the collimator angle $\vartheta = -\cos^{-1}(m_{1,1})$.

A.2. Standard Deviation of a Uniform Random Distribution

Consider the special case of the sequence s consisting of N equally spaced samples s = (1,2,3,...,N) with the constant (uniform) probability $p_i = \frac{1}{N}$ for the occurrence of any number between 1 and N. Then the mean of that distribution is obviously:

$$\bar{x}_N = \sum_{i=1}^{N} i \cdot p_i = \frac{1}{N}\sum_{i=1}^{N} i = \frac{1}{2}(N+1), \tag{A14}$$

and the variance σ^2 is:

$$\begin{aligned}\sigma_N^2 &= \sum_{i=1}^{N}(i-\bar{x})^2 \cdot p_i = \frac{1}{N}\sum_{i=1}^{N}\left(i^2 - 2i\bar{x} + \bar{x}^2\right) = \frac{1}{N}\left(\sum_{i=1}^{N} i^2 - 2\bar{x}\sum_{i=1}^{N} i + \bar{x}^2\sum_{i=1}^{N} 1\right) \\ &= \frac{1}{N}\left(\frac{N(N+1)(2N+1)}{6} - 2\frac{N}{2}(N+1)\cdot\frac{1}{2}(N+1) + N\left(\frac{1}{2}(N+1)\right)^2\right) \\ &= \frac{1}{12}\left(N^2-1\right)\end{aligned} \tag{A15}$$

If the sequence s is divided by N, this division carries through the calculation of $\bar{x}_N$ and σ^2, and σ^2 is obtained for N uniformly distributed random numbers in the interval [0, 1]:

$$\sigma^2 = \frac{\sigma_N^2}{N^2} = \frac{1}{12}\left(1 - \frac{1}{N^2}\right) . \tag{A16}$$

The standard deviation σ is defined as the square root of the variance, thus:

$$\sigma = \sqrt{\frac{1}{12}\left(1 - \frac{1}{N^2}\right)} . \tag{A17}$$

In the case of this thesis, there are 60 samples, given by 20 random angles of rotation about each of the three axes. Therefore N = 60 and σ = 0.29 for an interval of width 1. For an interval width of 2.5°, therefore σ = 0.72° is found.

A.3. References

1 G. Starkschall, "Coordinate systems in radiation oncology imaging", presented at the AAPM Summer School, Madison, Wi, (1998).

2 R. Onimaru, H. Shirato, H. Aoyama, K. Kitakura, T. Seki, K. Hida, K. Fujita, K. Kagei, T. Nishioka, T. Kunieda, Y. Iwasaki, and K. Miyasaka, "Calculation of rotational setup error using the real-time tracking radiation therapy (rtrt) system and its application to the treatment of spinal schwannoma", Int. J. Radiat. Oncol. Biol. Phys. **54**, 939-947 (2002).

Bibliography

(Page numbers in square brackets at the end of each entry indicate the location of each reference in this thesis)

Bansal, R., L. H. Staib, Z. Chen, A. Rangarajan, J. Knisely, R. Nath, and J. S. Duncan, "Entropy-based dual-portal-to-3-dct registration incorporating pixel correlation", IEEE Trans. Med. Imag. 22, 29-49 (2003). [p. 8]

Birkfellner, W., J. Wirth, W. Burgstaller, B. Baumann, H. Staedele, B. Hammer, N. C. Gellrich, A. L. Jacob, P. Regazzoni, and P. Messmer, "A faster method for 3d/2d medical image registration--a simulation study", Phys. Med. Biol. 48, 2665-2679 (2003). [p. 7]

de Boer, H. C., J. R. van Sornsen de Koste, C. L. Creutzberg, A. G. Visser, P. C. Levendag, and B. J. Heijmen, "Electronic portal image assisted reduction of systematic set-up errors in head and neck irradiation", Radiother. Onc. 61, 299-308 (2001). [p. 8]

Bortfeld, T., M. van Herk, and S. B. Jiang, "When should systematic patient positioning errors in radiotherapy be corrected?" Phys. Med. Biol. 47, N297-302 (2002). [p. 89]

Burman, C., G. J. Kutcher, B. Emami, and M. Goitein, "Fitting of normal tissue tolerance data to an analytic-function", International Journal of Radiation Oncology Biology Physics 21, 123-135 (1991). [p. 76]

Bushberg, J. T., J. A. Seibert, E. M. Leidholdt, and J. M. Boone, "The essential physics of medical imaging", 2nd ed, (Lippincott Williams & Wilkins, 2002). [p. 19]

Cho, B. C. J., M. van Herk, B. J. Mijnheer, and H. Bartelink, "The effect of set-up uncertainties, contour changes, and tissue inhomogeneities on target dose-volume histograms", Med. Phys. 29, 2305-2318 (2002). [p. 9]

Chung, P. W., T. Haycocks, T. Brown, Z. Cambridge, V. Kelly, H. Alasti, D. A. Jaffray, and C. N. Catton, "On-line asi portal imaging of implanted fiducial markers for the reduction of interfraction error during conformal radiotherapy of prostate carcinoma", Int. J. Radiat. Oncol. Biol. Phys. 60, 329-334 (2004). [p. 8]

Clippe, S., D. Sarrut, C. Malet, S. Miguet, C. Ginestet, and C. Carrie, "Patient setup error measurement using 3d intensity-based image registration techniques", Int. J. Radiat. Oncol. Biol. Phys. 56, 259-265 (2003). [p. 88]

Court, L., I. Rosen, R. Mohan, and L. Dong, "Evaluation of mechanical precision and alignment uncertainties for an integrated ct/linac system", Med. Phys. 30, 1198-1210 (2003). [p. 7]

Dekker, N., L. S. Ploeger, and M. van Herk, "Evaluation of cost functions for gray value matching of two-dimensional images in radiotherapy", Med. Phys. 30, 778-784 (2003). [p. 24]

Dong, L. and A. L. Boyer, "An image correlation procedure for digitally reconstructed radiographs and electronic portal images", Int. J. Radiat. Oncol. Biol. Phys. 33, 1053-1060 (1995). [p. 24]

v. Dyk, J., "The modern technology of radiation oncology", 1st ed., Vol. 1 (Medical Physics Publishing, 1999). [p. 22]

v. Dyk, J., "The modern technology of radiation oncology", 1st ed., Vol. 2 (Medical Physics Publishing, 2005). [p. 82]

Emami, B., J. Lyman, A. Brown, L. Coia, M. Goitein, J. E. Munzenrider, B. Shank, L. J. Solin, and M. Wesson, "Tolerance of normal tissue to therapeutic irradiation", International Journal of Radiation Oncology Biology Physics 21, 109-122 (1991). [p. 9]

Erridge, S. C., Y. Seppenwoolde, S. H. Muller, M. van Herk, K. De Jaeger, J. S. Belderbos, L. J. Boersma, and J. V. Lebesque, "Portal imaging to assess set-up errors, tumor motion and tumor shrinkage during conformal radiotherapy of non-small cell lung cancer", Radiother. Onc. 66, 75-85 (2003). [p. 8]

Fielding, A. L., P. M. Evans, and C. H. Clark, "Verification of patient position and delivery of imrt by electronic portal imaging", Radiother. Onc. 73, 339-347 (2004). [p. 8]

Fu, D., G. Kuduvalli, V. Mitrovic, W. Main, and T. Larry, "Automated skull tracking for the cyberknife (r) image-guided radiosurgery system", Medical Imaging 2005, Proc. of SPIE 5744, 1605-7422 (2005). [p. 7]

Gilhuijs, K. G., P. J. van de Ven, and M. van Herk, "Automatic three-dimensional inspection of patient setup in radiation therapy using portal images, simulator images, and computed tomography data", Med. Phys. 23, 389-399 (1996). [p. 8]

Gonzales R. C., and R. E. Woods, "Digital image processing", 2nd ed, (Prentice Hall, 2001). [p. 24]

Hajnal, J. V., D. L. G. Hill, and D. J. Hawkes, "Medical image registration", (CRC Press, 2001). [p. 24, 31]

Hanley, J., M. A. Lumley, G. S. Mageras, J. Sun, M. J. Zelefsky, S. A. Leibel, Z. Fuks, and G. J. Kutcher, "Measurement of patient positioning errors in three-dimensional conformal radiotherapy of the prostate", Int. J. Radiat. Oncol. Biol. Phys. 37, 435-444 (1997). [p. 6]

Happersett, L., G. S. Mageras, M. J. Zelefsky, C. M. Burman, S. A. Leibel, C. Chui, Z. Fuks, S. Bull, C. C. Ling, and G. J. Kutcher, "A study of the effects of internal organ motion on dose escalation in conformal prostate treatments", Radiother. Onc. 66, 263-270 (2003). [p. 83, 89]

vanHerk, M., A. Bruce, A. P. G. Kroes, T. Shouman, A. Touw, and J. V. Lebesque, "Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration", International Journal of Radiation Oncology Biology Physics 33, 1311-1320 (1995). [p. 83]

Herman, M. G. "Clinical use of electronic portal imaging", Semin Radiat Oncol 15, 157-167 (2005). [p. 8]

Hill, D. L. G., P. G. Batchelor, M. Holden et al., "Medical image registration", Phys. Med. Biol. 46, R1-R45 (2001). [p. 31]

Hoogeman, M. S., M. Van Herk, D. Yan, L. J. Boersma, P. C. M. Koper, and J. V. Lebesque, "A model to simulate day-to-day variations in rectum shape", International Journal of Radiation Oncology Biology Physics 54, 615-625 (2002). [p. 83]

Jans, H.-S., A. Syme, S. Rathee, and B. G. Fallone, "3d inter-fractional patient position verification using 2d-3d registration of orthogonal images", Med. Phys., in press. [p. 3]

Johns, H. E. and J. R. Cunningham, "The physics of radiology", 4th ed, (Charles C Thomas, 1983). [p. 20]

Jaffray, D. A. "Emergent technologies for 3-dimensional image-guided radiation delivery", Sem. Rad. Onc. 15, 208-216 (2005). [p. 7]

Jaffray, D. A., J. H. Siewerdsen, J. W. Wong, and A. A. Martinez, "Flat-panel cone-beam computed tomography for image-guided radiation therapy", International Journal of Radiation Oncology Biology Physics 53, 1337-1349 (2002). [p. 7]

Khan, F. M. "The physics of radiation therapy", 3 ed, (Lippincott Williams & Wilkins, 2003). [p. 4, 22]

Karger, C. P., D. Schulz-Ertner, B. H. Didinger, J. Debus, and O. Jakel, "Influence of setup errors on spinal cord dose and treatment plan quality for cervical spine tumours: A phantom study for photon imrt and heavy charged particle radiotherapy", Phys. Med. Biol. 48, 3171-3189 (2003). [p. 9]

Kitamura, K., H. Shirato, S. Shimizu, N. Shinohara, T. Harabayashi, T. Shimizu, Y. Kodama, H. Endo, R. Onimaru, S. Nishioka, H. Aoyama, K. Tsuchiya, and K. Miyasaka, "Registration accuracy and possible migration of internal fiducial gold marker implanted in prostate and liver treated with real-time tumor-tracking radiation therapy (rtrt)", Radiother. Onc. 62, 275-281 (2002). [p. 8]

Kress, J., S. Minohara, M. Endo, J. Debus, and T. Kanai, "Patient position verification using ct images", Med. Phys. 26, 941-948 (1999). [p. 7]

Kumar, S., K. Burke, C. Nalder, P. Jarrett, C. Mubata, R. A'Hern, M. Humphreys, M. Bidmead, and M. Brada, "Treatment accuracy of fractionated stereotactic radiotherapy", Radiother. Onc. 74, 53-59 (2005). 2[p. 8]

Lacroute, P. and M. Levoy, "Fast volume rendering using a shear-warp factorization of the viewing transformation", presented at the 21st Conference on Computer Graphics and Interactive Techniques, (1994). [p. 56]

Lemieux, L., R. Jagoe, D. R. Fish, N. D. Kitchen, and D. G. Thomas, "A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs", Med. Phys. 21, 1749-1760 (1994). [p. 7]

Lyman, J. T. "Complication probability as assessed from dose volume histograms", Radiation Research 104, S13-S19 (1985). [p. 9, 75]

Leszczynski, K. W., S. Loose, and S. Boyko, "An image registration scheme applied to verification of radiation therapy", Br. J. Radiol. 71, 413-426 (1998). [p. 8]

Matsopoulos, G. K., K. K. Delibasis, N. A. Mouravliansky, and K. S. Nikita, "Automatic medical image registration schemes using global optimization techniques", Stud Health Technol Inform 79, 463-491 (2000). [p. 8]

Murphy, M. J. "An automatic six-degree-of-freedom image registration algorithm for image-guided frameless stereotaxic radiosurgery", Med. Phys. 24, 857-866 (1997). [p. 7]

Ntasis, E., T. A. Maniatis, and K. S. Nikita, "Fourier volume rendering for real time preview of digital reconstructed radiographs: A web-based implementation", Comp. Med. Imag. Graph. 26, 1-8 (2002). [p. 58]

Okunieff, P., D. Morgan, A. Niemierko, and H. D. Suit, "Radiation dose-response of human tumors", International Journal of Radiation Oncology Biology Physics 32, 1227-1237 (1995). [p. 75]

Onimaru, R., H. Shirato, H. Aoyama, K. Kitakura, T. Seki, K. Hida, K. Fujita, K. Kagei, T. Nishioka, T. Kunieda, Y. Iwasaki, and K. Miyasaka, "Calculation of rotational setup error using the real-time tracking radiation therapy (rtrt) system and its application to the treatment of spinal schwannoma", Int. J. Radiat. Oncol. Biol. Phys. 54, 939-947 (2002). [p. 8, 97]

Penney, G. P., J. Weese, J. A. Little, P. Desmedt, D. L. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2-d-3-d medical image registration", IEEE Trans. Med. Imag. 17, 586-595 (1998). [p. 88]

Pisani, L., D. Lockman, D. Jaffray, D. Yan, A. Martinez, and J. Wong, "Setup error in radiotherapy: On-line correction using electronic kilovoltage and megavoltage radiographs", Int. J. Radiat. Oncol. Biol. Phys. 47, 825-839 (2000). [p. 7]

Pouliot, J., A. Bani-Hashemi, J. Chen, M. Svatos, F. Ghelmansarai, M. Mitschke, M. Aubin, P. Xia, O. Morin, K. Bucci, M. Roach, P. Hernandez, Z. R. Zheng, D. Hristov, and L. Verhey, "Low-dose megavoltage cone-beam ct for radiation therapy", International Journal of Radiation Oncology Biology Physics 61, 552-560 (2005). [p. 7]

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, "Numerical Recipes in C", 2nd ed, (Cambridge University Press, Cambridge New York Port Chester Melbourne Sydney, 1992). [p. 32]

Rehbinder, H., C. Forsgren, and J. Lof, "Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery", Med. Phys. 31, 3363-3371 (2004). [p. 89]

Remeijer, P., E. Geerlof, L. Ploeger, K. Gilhuijs, M. van Herk, and J. V. Lebesque, "3-d portal image analysis in clinical practice: An evaluation of 2-d and 3-d analysis techniques as applied to 30 prostate cancer patients", Int. J. Radiat. Oncol. Biol. Phys. 46, 1281-1290 (2000). [p. 6]

Rohlfing, T., D. B. Russakoff, J. Denzler, and C. R. Maurer, "Progressive attenuation fields: Fast 2d-3d image registration without precomputation", Medical Image Computing and Computer-Assisted Intervention - Miccai 2004, Pt 1, Proceedings 3216, 631-638 (2004). [p. 58]

Ryu, S., F. F. Yin, J. Rock, J. E. Zhu, A. Chu, E. Kagan, L. Rogers, M. Ajlouni, M. Rosenblum, and J. H. Kim, "Image-guided and intensity-modulated

radiosurgery for patients with spinal metastasis", Cancer 97, 2013-2018 (2003). [p. 7]

Sarrut, D. and S. Clippe, "Geometrical transformation approximation for 2d/3d intensity-based registration of portal images and ct scan", presented at the Lecture Notes In Computer Science, 4th MICCAI, (2001). [p. 57, 85]

Schaly, B., G. S. Bauman, W. Song, J. J. Battista, and J. Van Dyk, "Dosimetric impact of image-guided 3d conformal radiation therapy of prostate cancer", Phys. Med. Biol. 50, 3083-3101 (2005). [p. 83]

Shannon, C. E. "A mathematical theory of communication", Bell System Technical Journal 27, 379-423 (1948). [p. 31]

Sirois, L. M., D. H. Hristov, and B. G. Fallone, "Three-dimensional anatomy setup verification by correlation of orthogonal portal images and digitally reconstructed radiographs", Med. Phys. 26, 2422-2428 (1999). [p. 8, 35]

Soete, G., J. Van de Steene, D. Verellen, V. Vinh-Hung, D. Van den Berge, D. Michielsen, F. Keuppens, P. De Roover, and G. Storme, "Initial clinical experience with infrared-reflecting skin markers in the positioning of patients treated by conformal radiotherapy for prostate cancer", Int. J. Radiat. Oncol. Biol. Phys. 52, 694-698 (2002). [p. 7]

Smitsmans, M. H. P., J. de Bois, J. J. Sonke, A. Betgen, L. J. Zijp, D. A. Jaffray, J. V. Lebesque, and M. van Herk, "Automatic prostate localization on cone-beam ct scans for high precision image-guided radiotherapy", International Journal of Radiation Oncology Biology Physics 63, 975-984 (2005). [p. 7]

Song, W., B. Schaly, G. Bauman, J. Battista, and J. Van Dyk, "Image-guided adaptive radiation therapy (igart): Radiobiological and dose escalation considerations for localized carcinoma of the prostate", Med. Phys. 32, 2193-2203 (2005). [p. 9]

Starkschall, G. "Coordinate systems in radiation oncology imaging", presented at the AAPM Summer School, Madison, Wi, (1998). [p. 96]

Stroom, J. C., M. J. J. Olofsen-van Acht, S. Quint, M. Seven, M. De Hoog, C. L. Creutzberg, H. C. J. De Boer, and A. G. Visser, "On-line set-up corrections during radiotherapy of patients with gynecologic tumors", International Journal of Radiation Oncology Biology Physics 46, 499-506 (2000). [p. 6]

Van de Steene, J., F. Van den Heuvel, A. Bel, D. Verellen, J. De Mey, M. Noppen, M. De Beukeleer, and G. Storme, "Electronic portal imaging with on-line correction of setup error in thoracic irradiation: Clinical evaluation", International Journal of Radiation Oncology Biology Physics 40, 967-976 (1998). [p. 6]

Studholme, C., D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3d medical image alignment", Patt. Recog. 32, 71-86 (1999). [p. 88]

Thevenaz, P., U. E. Ruttimann, and M. Unser, "A pyramid approach to subpixel registration based on intensity", "IEEE Trans. Imag. Proc." 7, 27-41 (1998). [p. 19, 32]

Wang, F., T. E. Davis, and B. C. Vemuri, "Real-time drr generation using cylindrical harmonics", Lecture Notes in Computer Science 2489, 671-678 (2002). [p. 57]

Warkentin, B., P. Stavrev, N. Stavreva, C. Field, and B. G. Fallone, "A tcp-ntcp estimation module using dvhs and known radiobiological models and parameter sets", J. Appl. Clin. Med. Phys. 5, 50-63 (2004). [p. 76]

Weese, J., R. Goecke, G. P. Penney, P. Desmedt, T. M. Buzug, and H. Schumann, "Fast voxel-based 2d/3d registration algorithm using a volume rendering method based on the shear-warp factorization", presented at the SPIE Conf. Imag. Proc., San Diego, Ca, (1999). [p. 58]

Weese, J., T. M. Buzug, C. Lorenz, and C. Fassnacht, "An approach to 2d/3d registration of a vertebra in 2d x-ray fluoroscopies with 3d ct images", in Cvrmed-mrcas'97 (1997), Vol. 1205, pp. 119-128. [p. 87, 90]

Yan, D., B. Xu, D. Lockman, K. Kota, D. S. Brabbins, J. Wong, and A. A. Martinez, "The influence of interpatient and intrapatient rectum variation on external beam treatment of prostate cancer", International Journal of Radiation Oncology Biology Physics 51, 1111-1119 (2001). [p. 83]